



Institute for Scientific Computing Research

FY 1999 Annual Report

<http://www.llnl.gov/casc/iscr>

Lawrence Livermore National Laboratory
P.O. Box 808, L-561, Livermore, CA 94551





The University Relations Program (URP) encourages collaborative research between Lawrence Livermore National Laboratory (LLNL) and the University of California campuses. The Institute for Scientific Computing Research (ISCR) actively participates in such collaborative research, and this report details the Fiscal Year 1999 projects jointly served by URP and ISCR. For a full discussion of all URP projects in FY 1999, please request a copy of the URP FY 1999 Annual Report by contacting

Lawrence Livermore National Laboratory
Christina M. Budwine, University Relations Program
P. O. Box 808, L-413
Livermore, CA 94551

UCRL-LR-133866-99



DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from the Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (423) 576-8401

Available to the public from the National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd., Springfield, VA 22161

Contents

ISCR Fiscal Year 1999 Report	4
ISCR Fiscal Year 1999 in Review	6
Fiscal Year 1999 Seminar Series Abstracts	13
Visiting and Collaborating Professors' Project Abstracts.....	93
University Collaborative Research Program Subcontract Abstracts.....	103
Laboratory Directed Research and Development Project Abstracts	113
Student Internship Research Summaries.....	119
Scalable Linear Solvers Workshop Report	155

The Mission of the Institute for Scientific Computing Research

The Institute for Scientific Computing Research (ISCR) at Lawrence Livermore National Laboratory is jointly administered by the University Relations Program (URP) and the Center for Applied Scientific Computing (CASC), and this joint relationship expresses its mission. An extensively externally networked ISCR cost-effectively expands the level and scope of national computational science expertise available to the Laboratory, while streamlining the administrative burden that is unavoidable when bridging the Laboratory's internal computational research environment with that of the academic community.

As large-scale simulations on the parallel platforms of DOE's Accelerated Strategic Computing Initiative (ASCI) become increasingly important to the overall mission of LLNL, the role of the ISCR expands in importance, as well.

Relying primarily on nonpermanent staffing, the ISCR complements Laboratory research in areas of the computer and information sciences needed at the frontier of Laboratory missions. The ISCR works with CASC in being the "eyes and ears" of the Laboratory in the computer and information sciences, that is, in keeping the Laboratory aware of and connected to important advances. It also attempts to be "feet and hands" for the Laboratory by carrying those advances into the Laboratory and incorporating them into practice.

The ISCR has begun to and will increasingly provide continuing education opportunities for Laboratory personnel, in the form of on-site workshops taught by outside experts on novel software or hardware technologies.

The ISCR also seeks to influence the research community external to the Laboratory to pursue Laboratory-related interests, and to train the workforce that will be required by the Laboratory. Part of the performance of this function is interpreting to the external community appropriate (unclassified) aspects of the Laboratory's own contributions to the computer and information sciences — contributions that its unique mission and unique resources give it a unique opportunity and responsibility to make.

Of the three principal means of packaging scientific ideas for transfer — people, papers, and software — experience suggests that the most effective means is people. The programs of the ISCR are, therefore, people-intensive.

Finally, the ISCR, together with CASC, confers an organizational identity on the burgeoning computer and information sciences research activity at LLNL, and serves as a point of contact within the Laboratory for computer and information scientists from outside.

Institute for Scientific Computing Research

Fiscal Year 1999 Report

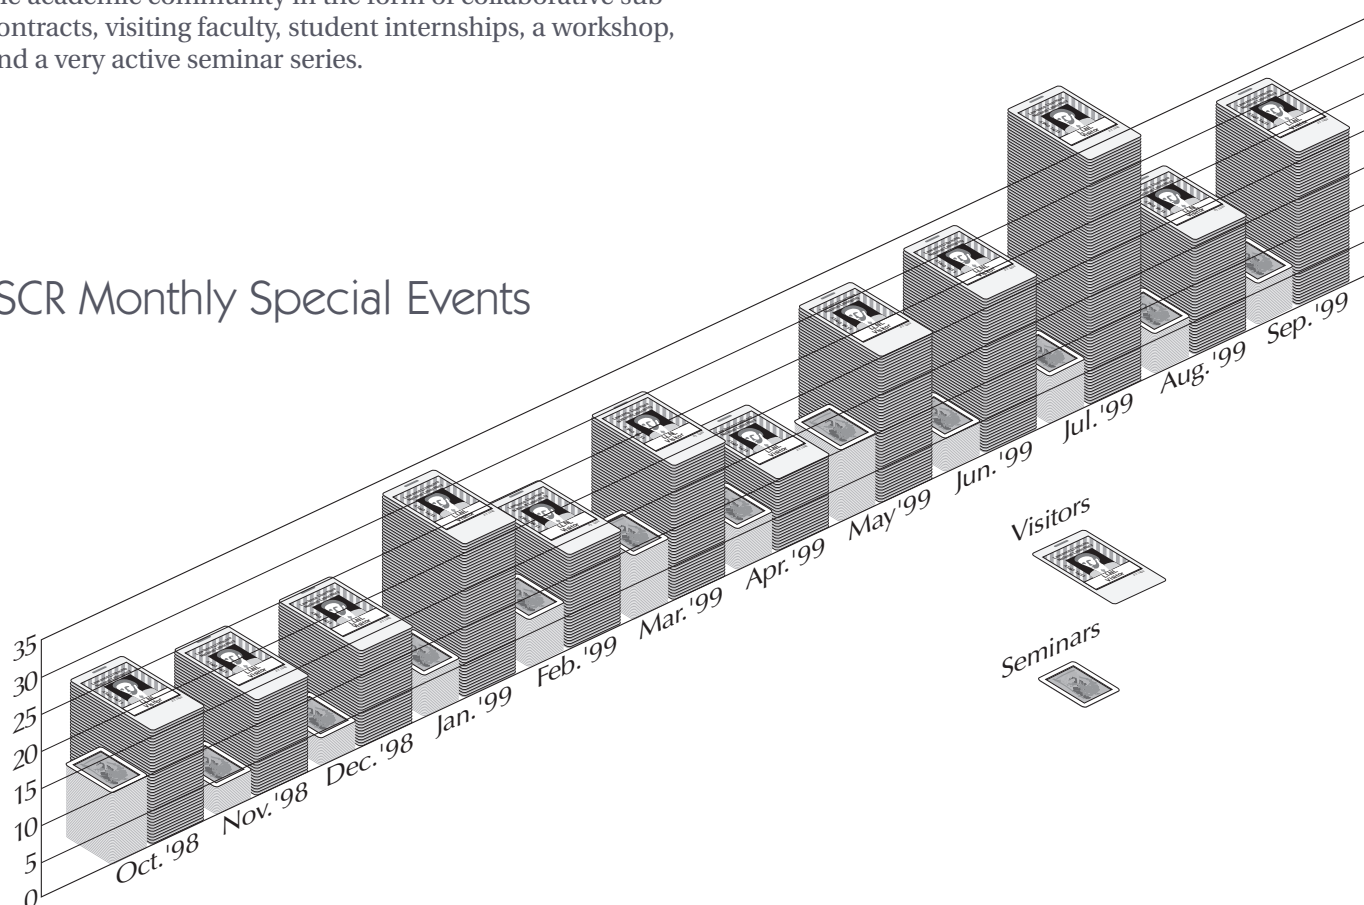
Large-scale scientific computation, and all of the disciplines that support it and help to validate it, have been placed at the focus of Lawrence Livermore National Laboratory by the Accelerated Strategic Computing Initiative (ASCI). The Laboratory operates the computer with the highest peak performance in the world and has undertaken some of the largest and most compute-intensive simulations ever performed. Computers at the architectural extremes, however, are notoriously difficult to use efficiently. Even such successes as the Laboratory's two Bell Prizes awarded in November 1999 only emphasize the need for much better ways of interacting with the results of large-scale simulations.

Advances in scientific computing research have, therefore, never been more vital to the core missions of the Laboratory than at present. Computational science is evolving so rapidly along every one of its research fronts that to remain on the leading edge, the Laboratory must engage researchers at many academic centers of excellence. In FY 1999, the Institute for Scientific Computing Research (ISCR) has expanded the Laboratory's bridge to the academic community in the form of collaborative sub-contracts, visiting faculty, student internships, a workshop, and a very active seminar series.

ISCR research participants are integrated almost seamlessly with the Laboratory's Center for Applied Scientific Computing (CASC), which, in turn, addresses computational challenges arising throughout the Laboratory. Administratively, the ISCR flourishes under the Laboratory's University Relations Program (URP). Together with the other four Institutes of the URP, it must navigate a course that allows the Laboratory to benefit from academic exchanges while preserving national security. Although FY 1999 brought more than its share of challenges to the operation of an academic-like research enterprise within the context of a national security laboratory, the results declare the challenges well met and well worth the continued effort.

A change of administration for the ISCR occurred during FY 1999. Acting Director John Fitzgerald retired from LLNL in August after 35 years of service, including the last two at helm of the ISCR. David Keyes, who has been a regular visitor in conjunction with ASCI scalable algorithms research since October 1997, overlapped with John for three months and serves half-time as the new Acting Director.

ISCR Monthly Special Events



The pages of this report summarize the activities of the faculty members, post-doctoral researchers, students, and guests from industry and other laboratories who participated in LLNL's computational mission under the auspices of the ISCR during FY 1999.

Altogether, the ISCR hosted 215 visits from 167 different visitors, who gave a total of 80 seminars on site. The vast majority of the visitors were from academia, with 14% from industry and 8% from other laboratories. Visitors from outside of the United States comprised 7% of the total. The histogram on the previous page shows visitors and seminars as a function of the month of the fiscal year. It displays year-round activity, with higher than average on-site collaborations in the summer.

Most of the material of this annual report comes directly from the visitors and principal investigators of the projects being reported, who selected formats convenient for their purposes. We thank Alane Alchorn and Dan Moore for collecting over a hundred separate pieces of text that make up this report and producing an easily navigated and visually pleasing document.

We hope that you enjoy examining this report on the ISCR's diverse activities in FY 1999. For further information about the Institute, please contact us at the address below.

Inquiries regarding the ways in which you might enhance the ISCR Programs in FY 2000 or beyond are welcome.


David Keyes



David E. Keyes
Acting Director, ISCR
925/ 422-1325
dekeyes@llnl.gov
<http://www.llnl.gov/casc/people/keyes>



Jill Dunaway
Administrator, ISCR
925/ 422-7132
dunaway4@llnl.gov
<http://www.llnl.gov/casc/people/dunaway>



Lawrence Livermore National Laboratory
P. O. Box 808, L-561
Livermore, CA 94551

<http://www.llnl.gov/casc/iscr>

ISCR Fiscal Year 1999 in Review

FY 1999 Seminar Series (in reverse chronological order)

David O'Hallaron, Carnegie Mellon University	September 17, 1999
Linda Petzold, University of California, Santa Barbara	September 13, 1999
Jarmo Rantakokko, University of California, San Diego	September 10, 1999
Gundolf Haase, Johannes Kepler University, Linz	September 7, 1999
Ulrich Langer, Johannes Kepler University, Linz	September 7, 1999
Balakrishna Iyer, IBM Santa Theresa Labs	September 3, 1999
Ioana Banicescu, Mississippi State University	August 26, 1999
Valerio Pascucci, University of Texas, Austin	August 25, 1999
Howard Elman, University of Maryland	August 24, 1999
Joerg Meyer, University of California, Davis	August 20, 1999
Gerik Scheuermann, University of California, Davis	August 20, 1999
Petra Stapf, Daimler Chrysler AG	August 17, 1999
Andreas Stathopoulos, College of William and Mary	July 29, 1999
Gerard Meurant, Commissariat à l'énergie atomique	July 26, 1999
Fernando Reitich, University of Minnesota	July 23, 1999
Krister Ahlander, Uppsala University	July 16, 1999
Raytcho Lazarov, Texas A&M University	July 12, 1999
Calton Pu, Oregon Graduate Institute	July 9, 1999
Robert Haimes, Massachusetts Institute of Technology	July 7, 1999
Alexander Schweitzer, Universitaet Bonn	June 18, 1999
Robert Grossman, University of Illinois at Chicago	June 15, 1999
Donald Schwendeman, Rensselaer Polytechnic Institute	June 14, 1999
Jonathan Shewchuk, University of California, Berkeley	June 9, 1999
Gene Golub, Stanford University	May 27, 1999
Michael Minion, University of North Carolina	May 26, 1999
Pieter Hemker, Centrum voor Wiskunde en Informatica	May 24, 1999
Paul Barton, Massachusetts Institute of Technology	May 20, 1999
Ethan Miller, University of Maryland, Baltimore County	May 18, 1999
Chaitanya Baru, University of California, San Diego	May 14, 1999
Barton Miller, University of Wisconsin	May 10, 1999
David Muraki, New York University	May 7, 1999
Craig Douglas, University of Kentucky	May 4, 1999
Ivan Yotov, University of Pittsburgh	May 3, 1999
Luc Vincent, Xerox Corporation	April 29, 1999
Matthew O'Keefe, University of Minnesota	April 28, 1999
Dennis Gannon, NASA Ames and Indiana University	April 27, 1999
Alan Heirich, Compaq Tandem Laboratories	April 20, 1999
Marsha Berger, New York University	April 14, 1999
Robert Haimes, Massachusetts Institute of Technology	April 1, 1999

Larry Snyder, University of Washington	March 25, 1999
Mei-Ling Liu, California Polytechnic State University	March 22, 1999
Peter Beckman, Los Alamos National Laboratory	March 17, 1999
Dinesh Manocha, University of North Carolina, Chapel Hill	March 16, 1999
Robert Lowrie, Los Alamos National Laboratory	March 5, 1999
Beth Wingate, Los Alamos National Laboratory	March 4, 1999
Andrea Malagoli, University of Chicago	March 4, 1999
Andy Wathen, Oxford University	March 3, 1999
Dimitri Mavriplis, NASA Langley Research Center	February 25, 1999
Anders Petersson, Chalmers University of Technology	February 23, 1999
Ronald DeVore, University of South Carolina	February 19, 1999
Paul Hovland, Argonne National Laboratory	February 10, 1999
James Bramble, Texas A&M University	February 9, 1999
Cheryl Fillekes, Doyle Bouzaid Sailmakers, New Zealand	February 8, 1999
Petri Fast, New York University	February 1, 1999
Yelena Yesha, University of Maryland and NASA Goddard	January 28, 1999
Lou Kondic, Duke University	January 25, 1999
Richard Ewing, Texas A&M University	January 14, 1999
Calvin Ribbens, Virginia Polytechnic Institute and State University	January 12, 1999
Joseph Pasciak, Texas A&M University	January 11, 1999
Hanan Samet, University of Maryland	January 8, 1999
Samuel Uselton, MRJ Technology Solutions	December 18, 1998
Heinz-Otto Kreiss, University of California, Los Angeles	December 15, 1998
Alex Pothén, Old Dominion University	December 11, 1998
Michael J. Holst, University of California, San Diego	December 10, 1998
David M. Cooper, Lawrence Livermore National Laboratory	November 24, 1998
Gail Carpenter, Boston University	November 6, 1998
George Karypis, University of Minnesota	November 5, 1998
Sergei Nepomnyaschikh, Russian Academy of Sciences	November 4, 1998
Roy Hemker, University of California, Los Angeles	October 29, 1998
Wray Buntine, Ultimode Systems, LLC and UC Berkeley	October 22, 1998
E. Tina Cheng, University of California, Davis	October 20, 1998
Ramdas Ram-Mohan, Worcester Polytechnic Institute	October 19, 1998
Roger Wets, University of California, Davis	October 16, 1998
Glen Niebur, University of California, Berkeley	October 13, 1998
Charles Hansen, University of Utah	October 12, 1998
Mark Adams, University of California, Berkeley	October 6, 1998
Justin Wan, University of California, Los Angeles	October 5, 1998
Elmer Lewis, Northwestern University	October 1, 1998

Visiting Faculty, Guests, Consultants, and Researchers

Visiting and Collaborating Professors

Jack Dongarra, University of Tennessee
Michael Holst, University of California, San Diego
David Keyes, Old Dominion University
Raytcho Lazarov, Texas A&M University
Cal Ribbens, Virginia Polytechnic Institute and State University
Jinchao Xu, Pennsylvania State University
Ytha Yu, California State University, Hayward

Participating Guests

Mark Adams, University of California, Berkeley
Marsha Berger, New York University
Marian Brezina, University of Colorado
George Byrne, Illinois Institute of Technology
David Dean, University of Colorado
John Fitzgerald, Lawrence Livermore National Laboratory (retired)
Kyle Gallivan, Florida State University
Bernd Hamann, University of California, Davis
Stanley Johnson, Lehigh University
Falko Kuester, University of California, Davis
Martin Lades, Gene Trace
David Larson, Bay Area Research Corporation
Andrea Malagoli, University of Chicago
Michael Minion, University of North Carolina
Joseph Pasciak, Texas A&M University
Elbridge Gerry Puckett, University of California, Davis
John Rice, University of California, Berkeley
Yousef Saad, University of Minnesota
Paul Saylor, University of Illinois
Jeffrey Scroggs, North Carolina State University
Daniel Wake, TMA, Inc.

Consultants

Nabil Adams, Rutgers University
Berni Alder, University of California, Professor Emeritus
Randolph Bank, University of California, San Diego
Leo Breiman, University of California, Berkeley
Harry Dwyer, University of California, Davis

Consultants (continued)

Anne Greenbaum, University of Washington
Chuck Hansen, University of Utah
David Keyes, Old Dominion University
Heinz-Otto Kreiss, University of California, Los Angeles
Thomas Manteuffel, University of Colorado
Stephen McCormick, University of Colorado
Linda Petzold, University of California, Santa Barbara
Steve Schaffer, New Mexico Institute of Technology
Homer Walker, Worcester Polytechnic Institute

Postdoctoral Researchers

Erick Cantú-Paz
Petri Fast
Raymond Fellers II
Imola Fodor
Barry Lee
Brian J. Miller
Thomas Rutaganira
Leonid Tsap

University Collaborative Research Program Subcontractors

Scott Baden, University of California, San Diego
Abdul Barakat, University of California, Davis
Jackson Beatty, University of California, Los Angeles
John Dawson, University of California, Los Angeles
Viktor Decyk, University of California, Los Angeles
Roy Hemker, University of California, Los Angeles
Warren Mori, University of California, Los Angeles
Linda Petzold, University of California, Santa Barbara
Fred Pollitz, University of California, Davis
Joachim Raeder, University of California, Los Angeles

LDRD Project Investigators

Peter Brown, LLNL Center for Applied Scientific Computing
Mark Duchaineau, LLNL Center for Applied Scientific Computing
Chandrika Kamath, LLNL Center for Applied Scientific Computing

Students and Faculty

Department of Applied Science Faculty

Nelson Max
Garry Rodrigue

Department of Applied Science Students

Henry J. Alme
Paul Covello
Rebecca Mattson Darlington
Joseph Koning
Daniel Laney
Sean Lehman
David Miller
Tim Pierce
Jonathan Rochez
Subhasis Saha
Jay Salmonson
Jay F. Thomas
Michael Everett Wickett

Visiting Students

Travis Austin, University of Colorado
Martin Bertram, University of California, Davis
Kathleen S. Bonnell, University of California, Davis
Thomas Brunner, University of Michigan
Timothy Chartier, University of Colorado
J. Kevin Durrenberger, California State University, Hayward
Daniel Ellis, Brigham Young University
Matt Giamporcaro, Boston University

Visiting Students (continued)

Ana Iontcheva, University of California, Davis
Mark Jeunnette, University of Chicago
Chisup Kim, Texas A&M University
Matthew Knepley, Purdue University
Eric LaMar, University of California, Davis
Gary Lee, University of California, Berkeley
Scott Morris, University of Utah
David Nault, University of Cincinnati
Jason Perry, University of Kentucky
Bobby Philip, University of Colorado
Alexandru Pomeranz, University of California, Davis
Lyn Reid, University of Washington
Kevin Scully, University of California, Irvine
Alex Tamasan, University of Washington
Heidi Thornquist, Rice University
Ben Tobin, Northern Arizona University
Stanimire Tomov, Texas A&M University
Kevin Vlack, University of Illinois, Urbana-Champaign
Christian Weiss, Technische Universitaet
Ludmil Zikatanov, Pennsylvania State University

National Physical Science Consortium (NPSC) Students

Lora Ballinger, University of Utah
Rachel Karchin, University of California, Santa Cruz
Imelda Kirby, University of Washington
Megan Thomas, University of California, Berkeley

Scalable Linear Solvers Workshop Leaders

Robert D. Falgout
Jinchao Xu



Institute for Scientific Computing Research

Fiscal Year 1999 Seminar Series Abstracts

(in reverse chronological order)



September 17, 1999

Dv: A Toolkit for Building Remote Visualization Services

David O'Hallaron

droh@cs.cmu.edu

Carnegie Mellon University

Abstract

The talk describes a toolkit being developed at Carnegie Mellon, called Dv, for building remote visualization services on computational grids. The long-range goal is develop a general solution to the problem of providing heavyweight Internet services—such as visualization, datamining and advanced search engines—that require significant computation in order to satisfy requests. Initially, Dv is being used to produce a grid-enabled version of vtk for visualizing the output of earthquake ground motion simulations produced by the CMU Quake project.

The Dv toolkit is based on a form of mobile object, called an active frame, that consists of Java and native application data, a Java application program that describes the visualization flowgraph, and a scheduler. Each node in the flowgraph corresponds to a vtk routine. Active frames are processed by active frame servers running on the grid. Each server waits for an active frame to arrive, invokes the scheduler to determine the portion of the flowgraph to execute and the identity of the next host, executes that part of the flowgraph using native vtk routines, constructs an output frame, and sends it to the next host.

A remote Dv visualization consists of a series of active frames that hop from server to server until they reach the client for display. The idea is to provide a flexible and uniform framework for partitioning applications between remote and local systems based on available computing and communication resources, and for studying and deploying application-level schedulers.

Speaker's web page: <http://www.cs.cmu.edu/~droh>

September 13, 1999

Algorithms and Software for Sensitivity Analysis and Optimal Control of Large-Scale Differential–Algebraic Systems

Linda Petzold

petzold@engineering.ucsb.edu

University of California, Santa Barbara

Abstract

In recent years, as computers and algorithms for simulation have become more efficient and reliable, an increasing amount of attention has focused on the more computationally intensive tasks of sensitivity analysis and optimal control. In this lecture we describe algorithms and software for sensitivity analysis and optimal control of large-scale differential–algebraic systems, focusing on the computational challenges. We introduce a new software package, DASPK3.0, for sensitivity analysis, and discuss our progress to date on the COOPT software and algorithms for optimal control which has been used successfully for a wide range of applications including spacecraft trajectory optimization, manufacture of superconducting thin films, and optimization of tissue engineering processes. Results for applications will be described as time permits.

Speaker's web page: <http://eci2.ucsb.edu/mee/petzold>

Institution web page: <http://www.ucsb.edu>

Friday, September 10, 1999

A Graph Partitioning Approach for Irregular Block Decomposition

Jarmo Rantakokko

jarmo@cs.ucsd.edu

University of California, San Diego

Abstract

Irregular block decompositions arise in parallel implementations of structured adaptive mesh refinement (SAMR) as well as in multiblock methods. In this talk we consider the problem of how to efficiently cover and partition an irregular domain arising in ocean modeling. The boundary, i.e., the coastline, is irregular and there may be disconnected islands within the domain. Our objective is to eliminate as many land points as possible without introducing high overhead costs, and to balance the workloads evenly. A similar problem also arises in SAMR where the flagged points cover regions of high error. We have developed a partitioning algorithm based on spectral graph decomposition, which includes tuning parameters for adjusting the granularity of partitioning, and procedures for fine-tuning the partitioning with respect to overhead costs and load imbalance.

We will discuss our algorithm in detail and compare it with two basic partitioning methods, commonly used in ocean modeling. Partitioning with our approach exhibit low communication overheads and load imbalance. We discuss performance results on ASCI Blue-Pacific and the IBM-SP2 and derive an analytic performance model which accurately predicts the running time of the application.

Institution web page: <http://www.ucsd.edu/>

September 7, 1999

Scientific Computing Tools for 3D Field Problems

Gundolf Haase and Ulrich Langer

ghaase@numa.uni-linz.ac.at

ulanger@numa.uni-linz.ac.at

Johannes Kepler University, Linz

Abstract

The 3D magnetic field problems are challenging not only for their interesting applications in industry, but also from the mathematical point of view. Usually, technical 3D magnetic field problems are characterized by complicated interface geometries with potentially moving parts (e.g. rotating parts), large coefficient jumps, non-linearities, singularities, and the necessity of calculating the exterior magnetic field. In practice, the aim of the simulation is often the optimization of the magnetic device we are dealing with. In order to handle such kind of technical magnetic field problems, it is not sufficient to have a fast Maxwell solver and an optimizer; a geometry modeler, an advanced 3D mesh generator, mesh-handling and refinement methods, parallelization tools, and postprocessing tools including advanced visualization techniques also are required.

We will present such pre- and postprocessing tools, specially adapted to the adaptive multilevel methods used in the solver and the optimizer. Of course, the heart of our magnetic field problem solver environment FEPP is an adaptive multilevel Maxwell solver. In the magnetostatic case, the Maxwell solver is based on special mixed variational formulations of the Maxwell equations and their discretization by the Nédélec and Lagrange finite elements. Combining this with an adaptive nested multilevel preconditioned iteration approach, we obtain an

optimal solver with respect to the complexity. This is confirmed by the results of numerical experiments for academic problems and real-life applications as well. We also propose a concept for coupling finite elements with boundary elements. Coupled finite and boundary element schemes are well suited for problems where it is necessary to take into account the exterior magnetic field.

For the parallelization, two different strategies have been developed. The first approach uses a thread-based implementation that is especially suited for shared memory parallel computers such as the ORIGIN 2000. It is highly efficient if small numbers of processors are used. The second concept is based on distributed data algorithms and has been developed for massively parallel computers and workstation clusters.

The algorithms developed are used for solving more challenging engineering problems including 3D transient magnetomechanical problems.

More information on the problem solver environment FEPP and on the presented results can be obtained from the homepage <http://www.sfb013.uni-linz.ac.at> of the special research project (SFB) "Numerical and Symbolic Scientific Computing" that is supported by the Austrian Science Foundation (FWF).

September 3, 1999

Intelligent Miner and the Data Mining Challenge

Balakrishna Iyer

balaiyer@us.ibm.com

IBM Santa Theresa Labs

Abstract

The decreasing cost of computation allows for significant amounts of computing to solve problems previously unaddressed or addressed by other means. The decreased cost of communications bandwidth, and standardization of communications protocols, have made it possible to transport vast amounts of remotely stored/generated data. The decreased cost of storage makes it possible to collect/store/replicate more data.

The defeat of the world's best chess player by a computer, although controversial, is a landmark in data analysis. Intelligent Miner is a toolkit to help data analysts. It has been used successfully to spot fraud in the health sector and to predict financial trends. The talk will describe the toolkit, the kernel algorithms in the toolkit, and the data transformation function it offers.

August 26, 1999

Recent Dynamic Scheduling Strategies for Scientific Computing

Ioana Banicescu

ioana@cs.msstate.edu

Mississippi State University

Abstract

Scientific problems are large, irregular, and result in computationally intensive applications. Although many scientific applications are amenable to parallel execution, high performance is difficult to obtain due to load imbalance, one of the essential performance degradation factors. Load imbalances result from problem, algorithmic, and systemic characteristics. Recently, a number of dynamic scheduling schemes have been proposed and implemented in scientific applications. Some of these schemes (factoring, fractiling, weighted factoring, factoring with adaptive weights) are based on a probabilistic analysis and thus, they accommodate load imbalances caused by predictable phenomena, such as irregular data, as well as unpredictable phenomena, such as data-access latency and operating system interference.

Factoring and fractiling were successfully implemented in Monte-Carlo simulations and N-Body simulations, respectively. Weighted factoring and factoring with

adaptive weights also proved to be cost-effective in computational field simulations using unstructured and adaptive grids. In this talk, I will report on our experiments with these techniques on IBM-SP2 and a SuperMSPARC. Performance of simulation codes was improved by as much as 52%. This improvement underscores the need for scheduling schemes that accommodate load imbalances due to application as well as system induced variances.

Future work will be dedicated to extensions of these methods, and to new techniques that could improve performance of scientific applications on parallel and distributed architectures.

Speaker's web page: <http://www.cs.msstate.edu/~ioana>

Research web page:
<http://www.erc.msstate.edu/~bilderba/research.html>

August 25, 1999

Scientific Visualization in ... Python?

Valerio Pascucci

pascucci@cs.utexas.edu

University of Texas, Austin

Abstract

Scientific visualization systems have a special need for the integration of many heterogeneous modules to perform numerical computations, large data management, graphics and client server communication among diverse platforms over the net.

A scripting language such as Python seem to encompass several interesting characteristics that are useful in the development of a scientific visualization system, especially in an experimental environment where continuous changes to several modules are developed and integrated independently. The missing element that I have been working on is a simple light-weighted visualization framework called "ndVP" in which the integration of independently developed modules becomes a trivial task. The system is designed to take full advantage of the Python flexibility, robustness and portability including the possibility of being modified/extended at runtime. Following the Python philosophy, the visualization framework proposed here involves a small kernel based on few basic classes: Groups, Models and

Display Windows. All the rest will grow as a set of external modules.

After an overview of the ndVP system I will present the details of a new isocontouring algorithm called Progressive Isocontouring that I have been developing as an extension module of ndVP. The Progressive Isocontouring is an isocontouring that incrementally builds a multiresolution representation of an isocontour. Progressive Isocontouring generates a data-structure that is consistent (valid mesh representation, no cracks, ...) at any time so that partial results can be displayed at any intermediate stage during its evaluation. In the ndVP system the Progressive Isocontouring algorithm has been tested using independent threads for contour evaluation and contour visualization. The visualization thread accesses the data-structure that the evaluation thread is building and performs the redraw action at regular intervals. If a request for a new isocontour is issued the main process terminates the isocontour computation thread and starts a new one.

August 24, 1999

A Multigrid Method Enhanced by Krylov Subspace Iteration for Discrete Helmholtz Equations

Howard Elman
elman@cs.umd.edu

University of Maryland

Abstract

Standard multigrid algorithms have proven ineffective for the solution of discretizations of Helmholtz equations. In this work we modify the standard algorithm by adding GMRES iterations at coarse levels and as an outer iteration. We demonstrate the algorithm's effectiveness through theoretical analysis of a model problem and experimental results. In particular, we show that the combined use of GMRES as a smoother and outer iteration produces an algorithm whose performance depends relatively mildly on wave number and is robust for normalized wave numbers as large as 200. For fixed wave numbers, it displays grid-independent convergence rates and has costs proportional to number of unknowns.

Speaker's web page: <http://www.cs.umd.edu/~elman/>

August 20, 1999

Interactive Visualization of Medical and Biological Data Sets

Joerg Meyer
jmeyer@ucdavis.edu

University of California, Davis

Abstract

Interactive visualization of large medical and biological data sets requires advanced techniques in image processing and data reduction. We have designed a platform-independent rendering system that allows the user to interact with an object on a screen with no delay, i.e., in the same way as he or she would handle an object in the real world. The system provides a user-friendly user interface and a pipeline architecture that maintains interactive behavior independent from the amount of data, which must be processed at a given time frame.

We describe an interactive rendering system that uses hierarchical, multiresolution representations of the data set, and scalable rendering algorithms. The system, called InVIS, is designed to always guarantee a certain frame rate by introducing a time-control unit.

The rendering pipeline accepts input data from CT or MRI scanners, confocal laser scanning microscopes, CAD systems, and other imaging techniques. A flexible import module makes it easy to customize the system and adapt it to user-specific needs. Due to hierarchical refinement, the system ensures interactive response at any time. The rendering stage is based on OpenGL, which makes it easily portable to different hardware platforms. Ongoing research focuses on large and time-variant data sets.

August 20, 1999

Topological Vector Field Visualization with Clifford Algebra

Gerik Scheuermann
scheuer@ucdavis.edu

University of California, Davis

Abstract

Topology is an important tool for the analysis of vector fields. In the talk, it will be shown that Clifford algebra can help in finding features such as non-linear critical points. Clifford algebra is a mathematical model extending the usual vector space description of geometry by a multiplication of vectors. A description will be given only to enable understanding the visualization results.

In a second part, the effect of a triangulation on the complexity of interpolated flows is analyzed, as are the ways in which this complexity can be reduced by simple changes in the grid connectivity. The complexity here is just the number of critical points in the flow. Its dependence on the triangulation is used to define a data-dependent triangulation that reduces this number by flipping edges. This can also be used for the triangulation of curvilinear grids.

The last part shows the topological results from using C1-interpolation as vector field model. It will be demonstrated that a higher order interpolation enhances the local topology extraction compared to linear methods.

August 17, 1999

Reduction of NO_x in Lean Exhaust by Selective NO_x-Recirculation (SNR-Technique) and an Outlook to the Progress in Soot Modeling for Engines

Petra Stapf

petra.stapf@daimlerchrysler.com

Daimler Chrysler AG

Abstract

The Selective NO_x-Recirculation Technique (SNR-Technique) is a new NO_x aftertreatment system for lean-burn gasoline and diesel applications. The objective of SNR is NO_x removal from lean exhaust gas by NO_x adsorption and subsequent selective external recirculation and decomposition of NO_x in the combustion process. The SNR project is composed of two major parts. First is the development of NO_x adsorbents that are able to store large quantities of NO_x in lean exhaust gas, and second is the NO_x decomposition by the combustion process.

A recent soot-formation model using detailed chemistry was modified for high pressure applications. Since the required computing times do not allow a direct implementation in engine codes, the soot model was incorporated into a dedicated group combustion model to assess its performance under high-pressure conditions. The combined model is capable of mimicking the microscopic details of combustion and pollutant formation in diesel engines.

July 29, 1999

A Parallel, Block, Jacobi–Davidson Implementation for Solving Large Eigenproblems on Coarse-Grain Environments

Andreas Stathopoulos
andreas@cs.wm.edu

College of William and Mary

Abstract

Iterative methods often provide the only means of solving large eigenvalue problems. Their block variants converge slowly but they are robust, particularly in the presence of multiplicities. Preconditioning is often used to improve convergence. Yet, for large matrices, the demands posed on the available computing resources are huge.

Clusters of workstations and SMPs are becoming the main computational tool for many scientific and engineering groups. In these environments, high communication costs suggest coarse grain implementations of iterative methods.

We have combined the block and preconditioning functionalities into a parallel implementation of a block Jacobi–Davidson method. We combine a fine-grain iterative scheme with the coarse grain capability that each processor can precondition different eigenpairs. We outline the design and present some timings and convergence results on a small workstation cluster and on a SUN Enterprise.

Speaker's web page: <http://www.cs.wm.edu/~andreas/>

July 26, 1999

The Computation of Bounds for the Norm of the Error in the Preconditioned Conjugate Gradient Method

Gerard Meurant

meurant@bruyeres cea.fr

Commissariat à l'énergie atomique

Abstract

In this talk we will first recall how to compute bounds for the A-norm of the error in the conjugate gradient (CG) method. This involves expressing the norm of the error as a Stieltjes integral. Then, quadrature formulas are used to approximate the integral. The Lanczos method is used to generate the nodes and weights of the quadrature rules. We will show how to use these techniques with CG and illustrate this by numerical experiments. Finally, we will extend these methods to the preconditioned version of CG and give numerical experiments showing that we can obtain very accurate upper and lower bounds for the norm of the error provided we have estimates of the smallest and largest eigenvalues of the matrix. This can be obtained adaptively during the first CG iterations leading to a reliable criterion for stopping PCG iterations.

July 23, 1999

Calculation of the Overall Magnetic Properties of Magnetorheological Fluids

Fernando Reitich
reitich@math.umn.edu

University of Minnesota

Abstract

Magnetorheological fluids (MRF), composed of micron-sized polarizable particles dispersed in a carrier liquid, constitute examples of controllable (“smart”) fluids whose rheological properties vary in response to an applied magnetic field. Understanding the magnetic behavior of MRF is crucial to the development of MRF-based devices and it also provides valuable insight into the character of the microstructure responsible for their field-dependent rheology. In this talk we will present results on the calculation of the overall magnetic response of MRF. We will show that effective medium approximations and particle dynamics simulations deliver numerical results that are in good agreement with experimental data and which can, therefore, be used to assist in the design of improved MRF.

Speaker's web page: <http://www.math.umn.edu/~reitich/>

Research web page:
<http://www4.ncsu.edu/~hvly/dynam.html>

July 16, 1999

COMPOSE: An Object-Oriented Framework for PDE Solvers

Krister Ahlander

krister@tdb.uu.se

Uppsala University

Abstract

The issue of how to develop reusable software in scientific computing is addressed. With object-oriented analysis and design an extendable set of collaborating objects, a framework named COMPOSE, has been developed for the numerical solution of partial differential equations.

Speaker's web page: <http://www.tdb.uu.se/~krister/>

Research web page:
<http://www.tdb.uu.se/research/swtools/>

July 12, 1999

Coupling Discretization Methods on Non-Matching Grids

Raytcho Lazarov
lazarov@rieman.llnl.gov

Texas A&M University

Abstract

The mortar method for coupling various approximation techniques has become an important tool in the construction and analysis of discretization schemes on non-matching grids.

First, we discuss briefly the main ingredients, issues, advantages, and the price we have to pay when using this approach. Second, on a differential level we introduce two hybrid formulations, which are the bases for the mortar and non-mortar approximations on non-matching grids.

Third, we discuss mortar finite volume methods and mortar mixed finite element discretizations and present the main steps in the stability and error analysis. Finally, we discuss iterative methods for solving the corresponding linear system and present some numerical experiments on model second-order problems.

July 9, 1999

An Extensible Wrapper Construction System for Internet Information Sources

Calton Pu

calton@cse.ogi.edu

Oregon Graduate Institute

Abstract

The amount of useful semi-structured data on the web continues to grow at a stunning pace. Often interesting web data are not in database systems but in HTML pages, XML pages, or text files. The data in these formats is not directly usable by standard SQL-like query processing engines that support sophisticated querying and reporting beyond keyword-based information retrieval. Hence, the web users or applications need a smart way of extracting data from these web sources. One of the popular approaches is to write wrappers around the sources, either manually or with software assistance, to bring the web data within the reach of more sophisticated query tools and general mediator-based information integration systems.

In this talk, we describe the methodology and the software development of a semi-automatic wrapper construction toolkit XWrap, which provides an XML-enabled, feedback-based, interactive wrapper construction facility

for generating value-added wrappers for Internet information sources. By XML-enabled, we mean that the extraction of information content from the web pages will be captured in an XML-compliant format and the query-based content filtering process is performed against XML documents. By feedback-based, we mean that the wrapper construction process is feedback-driven and the wrapper programs generated by XWrap can be incrementally revised and tuned according to the feedback collected through inductive learning by the wrapper generator. By value-added, we mean that any XWrap-generated wrapper program may accept more sophisticated queries than the corresponding web source can take. XWrap enables such enhanced functional capabilities through a smart composition of data wrapping and function wrapping. We will demonstrate the benefits of using wrappers generated by the XWrap in the context of the Continual Queries project.

Speaker's web page:

<http://www.cse.ogi.edu/~calton/>

July 7, 1999

Automated Feature Extraction from Transient CFD Simulations

Robert Haimes

haimes@orville.mit.edu

Massachusetts Institute of Technology

Abstract

In the past, feature extraction and identification were interesting concepts, but not required to fully understand the underlying physics of a steady-state flow field. This is because the results of the traditional visualization tools like iso-surfaces, cuts and streamlines are interactive and the results are easily represented to the user. These tools worked and properly conveyed the collected information at the expense of a great deal of interaction. For unsteady flow fields, the investigator does not have the luxury of spending time scanning only one “snap-shot” of the simulation. Automated assistance is required in pointing out areas of potential interest contained within the flow. This must not require a heavy compute burden (the extraction should not significantly slow down the solver for co-processing environments). With equal importance, methods must be developed to abstract the feature and display it in a manner that makes sense physically and is easily understood.

This talk presents the current status on locating shocks, vortices, separation surfaces (regions of recirculation) and boundary layers (and wakes) in steady-state and transient regimes.

June 18, 1999

A Particle-Partition of Unity Method for the Solution of Elliptic, Parabolic and Hyperbolic Partial Differential Equations

Alexander Schweitzer

schweitz@iam.uni-bonn.de

Universitaet Bonn

Abstract

In this paper, we present a meshless discretization technique for nonstationary convection–diffusion problems. It is based on operator splitting, the method of characteristics and a generalized partition of unity method. We focus on the discretization process and its quality. The method may be used as an h- or p-version. Even for general particle distributions, the convergence behavior of the different versions corresponds to that of the respective version of the finite element method on a uniform grid. We discuss the implementational aspects of the proposed method. Furthermore, we present the results of numerical examples, where we considered instationary convection-diffusion, instationary diffusion, linear advection and elliptic problems.

June 15, 1999

Tools and Techniques for High-Performance and Distributed Data Mining

Robert Grossman

grossman@uic.edu

University of Illinois at Chicago

Abstract

Data mining is the semi-automatic discovery of patterns, correlations, associations, changes and anomalies in large data sets. During the past two decades, the amount of data has grown explosively, yet the number of statisticians and other data analysts has remained relatively constant. There are only two possibilities: We either ignore larger and larger amounts of data, or we develop tools and techniques to aid in its automated analysis. In this talk, we give an overview data mining from this perspective and discuss some techniques for mining large and distributed data sets. We also provide an overview of some of the infrastructure that can make the mining of these types of data sets more practical.

June 14, 1999

An Adaptive Mesh Refinement Scheme for Hyperbolic Equations with Stiff Source Terms

Donald SchwendemanEmail: schwed@rpi.edu

Rensselaer Polytechnic Institute

Abstract

Multilevel adaptive mesh refinement (AMR) techniques have become popular for the numerical solution of non-linear hyperbolic equations in which sharp gradients and/or discontinuities (shock, detonations, etc.) in the solution develop and propagate. The basic idea is to refine the computational grid in regions where the solution changes rapidly in order to resolve these structures and to achieve high accuracy.

The talk will focus on current research aimed at the development of an AMR scheme for the numerical solution of hyperbolic equations with or without stiff source terms. The present scheme uses a second-order finite volume method on a structured hierarchical grid. The refinement is done by identifying grid patches where the solution changes rapidly in either space or time. Each

grid patch, in turn, can be refined so that the refinement scheme is recursive. Interpolation is used to generate fine grid patches from coarse grid data and to provide boundary data at fine-coarse grid boundaries. A description of the AMR scheme will be given for the case of unsteady, compressible, inviscid flow of a binary mixture of burnt and unburnt fuel as modeled by the reactive Euler equations.

Several examples will be used to illustrate the method and the special issues related to the numerical treatment of the stiff source term. Of particular interest will be the study of ignition and the evolution to detonation provoked by a small initial temperature gradient in an explosive sample, and the interaction of detonations created by several initial "hot spots."

June 9, 1999

Mesh Generation by Delaunay Refinement

Jonathan Shewchuk

jrs@cs.berkeley.edu

University of California, Berkeley

Abstract

Delaunay refinement is a technique for generating unstructured meshes of triangles or tetrahedra suitable for use in such applications as

- Graphics (especially radiosity).
- Terrain databases and Geographical Information Systems.
- Function interpolation.
- Numerical methods such as the finite element method.

In response to the needs of Carnegie Mellon's Quake project, an interdisciplinary Grand Challenge project whose goal is to simulate earthquake-induced ground motion in the Los Angeles basin, I have produced two general-purpose mesh generators. Triangle and Pyramid are implementations of two- and three-dimensional Delaunay refinement, respectively. Triangle has been released to the public, and has hundreds or thousands of users in both research and industrial settings, who use Triangle for all the applications listed above and more.

In this talk, I discuss the algorithmic underpinnings and current state of the art of Delaunay refinement methods. Popularized by the engineering community in the mid-1980s, Delaunay refinement operates by maintaining a Delaunay triangulation or Delaunay tetrahedralization, which is refined by the insertion of additional vertices. The placement of these vertices is chosen to force the mesh to conform to an object's boundaries, and to improve its quality as a grid for interpolation.

I present a new three-dimensional Delaunay refinement algorithm, which builds on the pioneering two-dimensional work by L. Paul Chew and Jim Ruppert. My algorithm produces tetrahedral meshes that are nicely graded (in a provable sense) and whose tetrahedra have circumradius-to-shortest edge ratios bounded below 1.63. This theoretical guarantee ensures that all poor quality tetrahedra except slivers (a particular type of poor tetrahedron) are removed. The slivers that remain are easily removed in practice, although there is no theoretical guarantee.

May 27, 1999

Iterative Methods for Solving Linear Systems

Gene Golub

golub@sccm.stanford.edu

Stanford University

Abstract

We discuss several problems in connection with solving linear systems. First, we consider various pre-conditioners for solving a variety of problems. This includes pre-conditioners for solving indefinite systems, and non-symmetric problems, especially when the skew-symmetric part of the matrix is dominant. Next, we describe the use of inner and outer iteration methods for solving problems where it is not possible to solve the equations using the pre-conditioner exactly. We attempt to restore the rate of convergence of the problem. Finally, we analyze the convergence properties of a method when the initial vector is considered to be a random vector. We give some numerical examples.

May 26, 1999

The Blob Projection Method for Immersed Boundary Problems

Michael Minion

minion@amath.unc.edu

University of North Carolina

Abstract

This talk will examine the problem of numerically approximating the evolution of a thin elastic membrane immersed in a constant density incompressible fluid. Immersed boundaries are currently used in modeling biological systems such as the beating of the heart and the swimming of flagellated organisms. A numerical method developed in collaboration with R. Cortez of Tulane will be presented which represents the membrane by a collection of regularized point forces while solving the incompressible Navier–Stokes equations on a regular Cartesian grid. The effects of the membrane on the fluid are calculated at grid points via a smoothed dipole potential and fast summation techniques, while the fluid equations are solved using a projection method. Comparisons between this method and the traditional immersed boundary method due to Peskin will also be presented.

May 24, 1999

Multigrid on 3D Adaptive Sparse Grids

Pieter Hemker
pieth@cw.nl

Centrum voor Wiskunde en Informatica

Abstract

We show some aspects of adaptive multigrid (MG) for the solution of 3D partial differential equations (PDEs). In particular, we discuss different possibilities of using partially ordered sets of auxiliary grids in multigrid algorithms. Because, for the classical sequence of grids, with 3D regular grids the number of degrees of freedom grows much faster with the refinement level than for 2D, it is more difficult to find sufficiently effective relaxation procedures and it makes sense to study the possibility of using larger sets of (regular rectangular) auxiliary grids.

Semi-coarsening is one technique in which a partially ordered set of grids is used. In this case still a unique fine-grid discrete problem is solved. On the other hand, sparse grid techniques may be more efficient if we compare the accuracy obtained with the number of degrees of freedom used. In the latter case, however, it is not straightforward to identify a discrete equation to be solved. Different approaches are compared.

We show different multigrid strategies, and results are given for transonic Euler-flow over the ONERA M6-wing and for a singular perturbation problem.

May 20, 1999

Theory, Numerics and Applications of Dynamic Sensitivity Analysis

Paul Barton
pib@mit.edu

Massachusetts Institute of Technology

Abstract

Dynamic sensitivity analysis has many applications in model reduction, model parameter estimation, controller design and numerical optimal control. The original existence and uniqueness results for dynamic sensitivities date back to Gronwall's work early this century. In this work, we extend Gronwall's results to include nonlinear ODEs with general consistent initial conditions, and linear time invariant DAEs. Furthermore, we develop a framework, and existence and uniqueness results, for sensitivity analysis of hybrid discrete/continuous dynamic systems with these differential systems embedded. Highlights of these analyses include the qualitative "jumping" of sensitivities at implicit discontinuities, and nonexistence of sensitivities at critical parameter values related to qualitative changes in the solution of the original dynamic system.

From the numerical standpoint, we discuss a novel and efficient algorithm for computing sensitivities simultaneously with BDF integration of the ODEs/DAEs. This algorithm is then extended to compute the sensitivities of hybrid systems. We present ABACUSS, a high level equation-based modeling environment for simulation and sensitivity analysis. The above theoretical results have some important practical implications for the development of robust general-purpose sensitivity software.

Finally, we discuss some practical applications of sensitivity analysis using ABACUSS, including model development and reduction for a citric acid process, and numerical optimal control. We conclude with some speculation concerning the role of sensitivity analysis in numerical technologies for the global solution of mixed-integer dynamic optimization problems.

May 18, 1999

Building Really Huge Scalable Storage Systems

Ethan Miller

elm@csee.umbc.edu

University of Maryland, Baltimore County

Abstract

In recent years, we have seen rapid increases in processor speed, memory capacity, and storage capacity. However, file system bandwidth has not kept up because of various limitations, both hardware and software. This talk will describe the storage challenges posed by slow hardware and inefficient hardware, and the research in parallel file systems and mass storage that we are doing to address these issues.

The first part of the talk will describe RAMA, our massively parallel file system. We are building a prototype whose bandwidth goal is 1 GB/sec using 100–200 commodity disks on commodity PC's connected by a Fast Ethernet switch. We can achieve this bandwidth by eliminating all centralized operations on file reads and writes, and by guaranteeing file system consistency even when file metadata is not written out immediately, reducing the number of disk I/Os per file block written. In addition to providing high bandwidth, this file system will parity-pro-

tect individual files, allowing a per-file tradeoff between performance and protection against disk failure. Moreover, this file system can accommodate large numbers of small files along side gigabyte (and larger) files without reducing performance on large files. We are currently testing a small prototype of this file system, and hope to deploy a larger prototype over the summer.

The second part of the talk will discuss the issues facing builders of petabyte (and larger) file systems. In particular, designers of such systems must deal with bandwidth that will inevitably lag behind increases in storage density. This problem affects access bandwidth, file system backup, and migration to newer storage technologies. I will quantify the problem and describe possible approaches to addressing it. Because this part of the talk is more speculative, feedback on both issues and directions will be welcome.

May 14, 1999

Integrating Archival Storage into Information Management Systems

Chaitanya Baru
baru@SDSC.EDU

University of California, San Diego

Abstract

Organizations are increasingly having to deal with archiving of large amounts of data over long periods of time. This has been a well known requirement in computational science and supercomputing applications where much effort is expended in creating large data collections (e.g. observational data, simulation outputs), and where it is typical for many applications to routinely generate large outputs.

We will discuss issues in combining the functionalities of archival storage systems with database management systems to efficiently support search and management of very large data collections. Archival storage systems such as the High Performance Storage System (HPSS) are designed to efficiently manage relatively small numbers (e.g. millions) of large data sets (e.g. hundreds of MB and higher). Combining DBMS technology with archival systems can allow efficient management of collections con-

taining large numbers (e.g. tens to hundreds of millions) of relatively small data sets (e.g. a few KB, or less, each). An example of such a system is the DB2/HPSS project, which is a joint effort between the IBM T.J.Watson Research Center and San Diego Supercomputer Center (SDSC), where the DB2 UDB system has been “integrated” with HPSS. We will provide details of this system and also describe other approaches to handling the “small data problem” in HPSS, including the “container” concept being developed for the Storage Resource Broker middleware at SDSC.

Finally, we will describe some projects at SDSC that are employing the above technologies. These include a project with the USPTO where the data collection is 2 million US patents, and a project with the California Digital Library (CDL), where the collection is about 100,000 art images from the Art Museum Image Consortium (AMICO).

May 10, 1999

Adaptive Operating Systems: An Architecture for Evolving Systems

Barton Miller

bart@cs.wisc.edu

University of Wisconsin

Abstract

Operating systems formerly were viewed as static entities, changing almost as slowly as the underlying hardware. Recent commercial systems have provided for a small amount of run-time change by allowing device drivers to be installed while the system is running. We are developing “adaptive operating systems” whose code can change and evolve while the system is running. This adaptation can be used to instrument the code for profiling or debugging purposes, or to modify and extend the operating system to adapt to changing work loads, application-demands, and configuration.

Our work differs from other efforts in this area in two ways. First, we can instrument and modify a stock, commercial operating system (Solaris) as it was delivered to the customer. We operate directly on the executable code while it is running. Second, we can modify the operating system at almost any point in its code. We are not constrained to system call or procedure call replacements.

Our research is embodied in a facility called KernInst. We will describe the basic KernInst mechanism and several uses of KernInst, including performance profiling and dynamically modifying and customizing the operating system in response to its work load. We will also present a case study using KernInst to profile the Solaris kernel under a Web proxy server (Squid) workload.

May 7, 1999

Lows, Highs, Jets & Storms: A Mathematical View of the Weather

David Muraki

muraki@cims.nyu.edu

New York University

Abstract

The most prominent features of the North American weather pattern, as often highlighted in TV weather segments, are the (west-to-east) jetstream and pressure systems of relative lows and highs. These behaviors are also reproduced in computations of simple weather models. The mathematics of the atmosphere is really a problem in fluid mechanics, but specialized to the situation where density/temperature effects and Coriolis force (due to the Earth's rotation) play dominant roles.

The basic equations for the midlatitude atmosphere are best understood in the limit of small Rossby number, an important theory known as “quasi-geostrophy.” As a leading-order asymptotic theory, quasi-geostrophy captures much of the dynamics of the atmosphere, including the fundamental instability of the jetstream

to pressure cells (falling pressure means bad weather). This theory, however, fails to explain the preferential development of low-pressure cells that lead to the asymmetric spirals characteristic of baroclinic storm systems.

By resolving a degeneracy inherent in this small Rossby number limit, a systematic asymptotic extension to the theory of quasi-geostrophy has been developed. Our computations of the next-order corrections are in excellent agreement with simulations of the full model equations, and provide a basis for explaining the asymmetric development of storms. Applications to other atmospheric phenomena, in collaboration with meteorologists at the National Center for Atmospheric Research (NCAR Boulder), are also discussed.

May 4, 1999

Cache-Based Multigrid Algorithms

Craig Douglas

douglas@ccs.uky.edu

University of Kentucky

Abstract

Multigrid methods combine a number of standard sparse matrix techniques. Usual implementations separate the individual components (e.g., an iterative method, residual computation, and interpolation between grids) into nicely structured routines. Today, however, many computers employ quite sophisticated and potentially large caches whose correct use is instrumental in gaining much of the peak performance of the processors. This is true independent of how many processors are used in a computation.

We investigate when it makes sense to combine several of the multigrid components into one, using block-oriented algorithms. We determine how large (or small) the blocks must be in order for the data in the block to just fit into the processor's primary cache. By re-using the data in cache several times, a potential savings in run time can be pre-

dicted. This is analyzed for a set of examples. It is surprising to see how large a subdomain can fit into a relatively small, well designed cache (e.g., 256KB). As caches continue to increase in size, the ideas here will extend quite nicely to three-dimensional problems. In particular, machines with 1–16MB caches are under design which will make three-dimensional cache-based multigrid a reality.

While most of the savings in time are with respect to the approximate solver, using a multiplicative or additive domain decomposition method saves even more time and allows us to use theoretical convergence rates from that field for our problems.

An automatic, low overhead software tool for determining a good cache based ordering for the multigrid components for general problems will also be discussed.

May 3, 1999

On the Accuracy and Efficiency of Coupling Different Grids and Discretizations in Flow in Porous Media Simulations

Ivan Yotov

yotov@math.pitt.edu

University of Pittsburgh

Abstract

Combining different grids and discretizations in a single numerical simulation has gained much popularity in recent years. It provides means for efficient modeling of irregular geometries, internal features, and local solution behavior. Critical for the success of this approach is imposing physically meaningful interface matching conditions in a numerically stable and accurate way. In this talk we consider flow in multiblock porous media domains discretized by mixed finite element methods on non-matching grids. We propose several mortar schemes (using specially chosen interface finite elements) and a non-mortar scheme for coupling along interfaces. The accuracy and efficiency of these approaches are compared theoretically and numerically. Efficient parallel implementation using a non-overlapping domain decomposition is discussed. The algorithms are illustrated by numerical simulations of multiphase flow in irregular heterogeneous domains.

April 29, 1999

Recent Advances in Morphological Methods for Image Segmentation and Feature Extraction

Luc Vincent

lucv@adoc.xerox.com

Xerox Corporation

Abstract

Mathematical morphology is an image processing methodology that was born in France in the late Sixties and has been gaining increasing importance and popularity since then. Today, it is broadly accepted as a powerful alternative to traditional linear techniques. Its rich set of operators is especially well suited for complex segmentation and feature extraction problems. Morphology has been successfully applied in many areas, including medical imaging, biology, radar, sonar, infra-red, and remote sensing images, industrial inspection, material science, fingerprints, identification, document recognition, etc.

The morphological approach to image analysis is natural and attractive: Binary images are considered as sets, whereas grayscale images are viewed as functions or topographic reliefs. These sets and functions are then transformed—in the spatial domain—via morphological operators, whose definitions are usually based on structuring elements, i.e., particular shapes that are translated in images and used as probes.

The presentation will start with a quick introduction to morphology's most well-known operations, namely erosions, dilations, openings, and closings, which are quite

useful for simple noise filtering and feature extraction tasks. We then move on to “granulometries.” These operations, based on openings and closings, are particularly interesting for extracting size information in images, for feature extraction, or for texture characterization. In addition, recent algorithmic advances have cut down their computation time by orders of magnitude, thereby opening up a range of new potential applications for these operators. Finally, we focus on segmentation. The watershed transformation forms the cornerstone of the morphological approach to segmentation. By constraining watersheds using previously extracted object markers (or “seeds”), a particularly powerful segmentation methodology is derived. Analyzing cases where watersheds do not perform well, However, leads us to propose alternative methods based on gray-values distance transforms and minimal paths. Among others, these methods are very efficient for the extraction of faint linear structure in grayscale images.

Though the talk covers a broad range of topics, it does not require any particular background in morphology from attendees. Operations will be described in an intuitive manner, and a large number of examples of applications will be shown.

April 28, 1999

GFS: A New File System Architecture for Network Devices

Matthew O'Keefe

okeefe@ee.umn.edu

University of Minnesota

Abstract

In computer systems today, speed and responsiveness are determined by network and storage subsystem performance. Faster, more scalable networking interfaces such as Fibre Channel and Gigabit Ethernet provide the scaffolding upon which higher-performance implementations may be constructed, but new thinking is required about how machines interact with smart network and storage devices

To this end, the Global File System (GFS) group's goal is to develop GFS, a scalable server-less file system that integrates IP-based network attached storage (NAS) and Fibre-Channel-based storage area networks (SAN). We call this new architecture Storage Area InterNeTworking (SAINT). It exploits the speed and device scalability of SAN clusters, and provides the client with the scalability and network interoperability of NAS appliances. GFS works under Linux and IRIX, though 95% of our efforts are now focused on Linux. GFS is freely-available, open source code released under the GNU General Public License. In this talk I will describe how GFS can provide the underlying parallel file system in large Linux clusters.

More information on GFS can be found at
<http://gfs.lcse.umn.edu>.

April 27, 1999

Java Grande: Prospects for High-Performance Computing Using Java

Dennis Gannon
gannon@cs.indiana.edu

NASA Ames and Indiana University,
Dept. of Computer Science

Abstract

This talk will present an overview of the Java Grande Forum and its valiant efforts to make Java safe for high-performance computing. Java Grande consists of two working groups. The forum is led by Geoffrey Fox. One working group is focused on the issue of numerics and the associated language features deemed necessary to make it a productive tool for large scale scientific computation. The other working group is focused around the area of parallelism and concurrency. This subgroup is working on ways to improve the efficiency of Java serialization and remote method invocation. Members are also working on proposals for Java MPI, though the talk will not dwell on this topic.

April 20, 1999

Alpha-Linux-ServerNet HPC Clusters and Visualization

Alan Heirich

alan.heirich@compaq.com

Compaq Tandem Laboratories

Abstract

This talk will present a strategy being executed by the Tandem division and Custom Systems Section of Compaq to develop a “Beowulf-class” cluster computing technology, and a related effort in scalable visualization. The clusters are built using the ServerNet-II SAN, a commodity HPC network technology with 350 MB/s bandwidth and application-to-application software latencies below 10 microseconds.

As part of this effort we are developing a scalable image combining architecture using 100% COTS components. The architecture computes the pixel-by-pixel merge of a set of images. It achieves low latency and high frame rates by using data compression and aggressive pipelining. The architecture uses gigabit crossbar switching in a fat tree topology to implement operator associativity, and uses an in-band control protocol to support high-speed mode changes in multi-pass rendering algorithms. The pixel combining function is user-specifiable and may involve pixel components (r,g,b,z,alpha,stencil) both as operands and as results. One use of the architecture is for interactive visualization of large-scale problems in computational science running on Beowulf-class workstation clusters. These clusters may contain dozens or hundreds of computational nodes and this architecture is designed to achieve interactive latency at these scales. This talk will discuss some requirements of algorithms for volume and vector field visualization and for photo-realistic lighting calculations.

April 14, 1999

Automatic High-Performance Fluid Computations in Complex Geometry

Marsha Berger

berger@nyu.edu

New York University

Abstract

We review aspects of our work on efficient and robust computational methods for computing moderate to high Mach number flows in complicated geometries. Our approach uses regular Cartesian meshes where the geometry intersects the mesh in an essentially arbitrary way. We draw on a variety of techniques in developing this approach. Adaptive mesh refinement is used to concentrate the computational effort in regions requiring greater resolution. The fluid flow equations need to be discretized in a special way to handle the cut cells where the geometry intersects the grid. Efficient algorithms from computational geometry allow us to find the cut cells and set up the discretization directly from a triangulation of the surface. We have taken great care to structure the geometric and flow computations so as to make the best use of cache-based moderately parallel computers.

April 1, 1999

pV3 — parallel Visual3k

Robert Haimes

haimes@orville.mit.edu

Massachusetts Institute of Technology

Abstract

pV3 builds heavily on the technology developed for Visual3. Visual3 is an interactive environment for visualization of 3D structured or unstructured data, on workstations with graphics hardware. The data may be steady or time-varying. Using Visual3, one can examine and probe a computational grid, as well as, display isosurfaces, cutting planes, and number of vector visualization functions such as streamlines and ribbons. Visual3 was originally developed for CFD applications.

pV3 is specifically designed for visualization of distributed data using a coprocessing paradigm. To create the visualization, pV3 runs simultaneously on a graphics workstation and an MPP, communicating by pvm. The workstation controls and displays the visualization. On the MPP, pV3 is coupled with the simulation and extracts the desired data for visualization and sends it to the workstation. pV3 has been designed to minimize network traffic. The client-side library extracts lower dimensional data required by the requested visualization tool from the volume of data in place.

March 25, 1999

Experimenting with ZPL

Larry Snyder

snyder@cs.washington.edu

University of Washington

Abstract

ZPL is a recently developed programming language designed to be fast, portable and convenient for scientific and engineering computations. ZPL has achieved these goals — experimental evidence will be presented in the lecture — because it was developed from first principles. Specifics of ZPL's machine and performance models will also be covered. The message is that programmers can write efficient, machine independent parallel programs without resorting to message passing.

ZPL was released in July 1997 and has been in regular use by a modest, but growing user community. Some users are simply solving their research problems. Others are experimenting with it because ZPL is the data parallel subset of a fully general language, Advanced ZPL. They want to become familiar with ZPL's approach in order to influence A-ZPL's design. The lecture will describe the evolutionary path to Advanced ZPL, with emphasis on its direct support for sparse arrays.

March 22, 1999

Caching Dynamic Web Objects and Measuring Network Performance

Mei-Ling Liu
mliu@csc.calpoly.edu

California Polytechnic State University

Abstract

In this talk, the presenter describes two of her projects at Cal Poly where, she conducts research with the assistance of Master's thesis students.

The first project studies the known techniques for reducing the latency on the web by using "Optimistic Delta", an approach developed by Lucent Technologies. A version of the approach is being implemented on the Cal Poly systems, and the performance impact of the technique is being measured. Our hope is to come up with improvements on the existing techniques.

In another project, a group of students performs measurements on the latency, throughput, and CPU utilization of the TCP/IP protocol stack on the Microsoft NT platform. The goal of the project is to develop a model to characterize the behavior of the protocol stack, in an attempt to maximize the network throughput on NT systems.

March 17, 1999

A Investigation of Myrinet, Linux, and Cost for High-Performance Scientific Computation

Peter Beckman

beckman@lanl.gov

Los Alamos National Laboratory

Abstract

Linux, the free, Open Source operating system, has seen explosive growth the last couple of years. According to a study from International Data Corporation, in the past year, the market share of Linux leapt from 6.8% to an estimated 17.2% of server operating system shipments. At Los Alamos National Laboratory, Linux has made similar inroads. Linux is being used lab-wide for desktop workstations, software development platforms, file and computation servers, and web and email servers.

Furthermore, Linux has proved to be extremely versatile for high-performance parallel computation on clusters of workstations. In 1994, NASA pioneered the use of Linux for building extremely cheap clusters with the Beowulf project. Since that time, dozens of laboratories, universities, and companies have been exploring Linux for high-performance computation. At Los Alamos, we have been exploring the performance and stability of our "Little Blue Penguin" cluster, a 64 dual-node (128 CPUs) Myrinet-connected parallel computer, across several different scientific applications. We have also looked at the cost effectiveness of high-speed interconnects for Linux clusters, and the current software available for parallel scientific programming.

March 16, 1999

A Framework for the Real-Time Walkthrough of Massive Models

Dinesh Manocha

dm@cs.unc.edu

University of North Carolina, Chapel Hill

Abstract

Computer-aided design (CAD) applications and scientific visualizations often need user-steered interactive displays or walkthroughs of very complex environments. Structural and mechanical designers often create models of ships, oil platforms, spacecraft, process plants and urban environments whose complexity exceeds the interactive visualization capabilities of current graphics systems. Yet for such structures the design process (and especially the multi-disciplinary design review process) benefits greatly from interactive walkthroughs.

In this talk we give an overview of our recent work on interactive walkthrough of massive models. These include new algorithmic approaches for accelerating rendering based on visibility culling, model simplification and image-based representations. We also present a framework for rendering very large models at nearly interactive rates. The framework scales with model size. Our framework can integrate multiple-rendering acceleration techniques, including visibility culling, geometric levels-of-detail and image-based approaches. We describe the database representation scheme for massive models used by the framework and a pipeline to manage the allocation of system resources among different techniques. We demonstrate the system of a coal-fired power plant composed of more than 15 million triangles.

March 5, 1999

Methods for Hyperbolic Systems with Stiff Relaxation

Robert Lowrie

lowrie@lanl.gov

Los Alamos National Laboratory

Abstract

A major challenge for numerical methods is the solution of hyperbolic systems with stiff relaxation terms. Applications include chemically reacting flows, multi-phase flows, elastic-plastic solid dynamics, gas kinetics, and radiation transport. It is usually straightforward to develop methods that are accurate when all of the time and length scales are resolved. When the time scale of the relaxation is unresolved, however, it is well known that many popular methods fail in that they are inconsistent with the correct asymptotic behavior (the Chapman-Enskog, continuum, or diffusion limit). In applications such as radiation hydrodynamics, it is critical that the correct limit be represented, including an accurate resolution of the diffusion term.

In this talk I will focus on shock-capturing, high-resolution ("Godunov-based") schemes. Even if the relaxation term is treated implicitly, unless an intricate time-stepping scheme is used (such as that due to Jin), conventional finite-volume methods have the wrong asymptotic behavior. However, using a simple time-integration method, the discontinuous Galerkin method obtains the correct diffusion limit for many problems. Results will be shown for radiation hydrodynamics, a simple elastic-plastic model, and the Broadwell model of gas kinetics. I will also discuss the requisite modification of the flux solver.

March 4, 1999

High-Order Spectral Elements on Triangles and Other Domains

Beth Wingate
wingate@lanl.gov

Los Alamos National Laboratory

Abstract

In this work we explore high-order spectral finite elements on domains such as triangles and tetrahedra. In the quadrilateral and cube, both the space (the tensor products of Legendre polynomials), and the points (the Gauss-Lobatto points), are well known. In the simplex, however, things are not as well understood. Current methods often give ill-conditioned matrices. Preconditioning can fix some of this at moderate order, but at higher-degree this fails. Part of the problem lies in the choice of space, and its basis. The other difficulty lies in the choice of points. For example, it doesn't appear that there exists a general formula for optimal (same number of points as degree of space) Gaussian quadrature in the triangle.

We present the multi-dimensional extension of the Legendre (and Tchebyshev) polynomials over the simplex. In the triangle these are the Koornwinder polynomials (1974). We are also able to compute points which give well-behaved Lagrange polynomials for the triangle and tetrahedra, the Fekete points, which are analogous in some ways to the Gauss-Lobatto points. These points also give a Gauss-like quadrature which has desirable symmetry properties.

This work has applications not only to the spectral element method, but to accurate interpolation in general domains, such as on the sphere, and to accurate and symmetric integration over the triangle and tetrahedra.

March 4, 1999

Simulations of Compressible MHD Problems Using Higher-Order Godunov Methods

Andrea Malagoli

a-malagoli@uchicago.edu

University of Chicago

Abstract

Higher-order Godunov methods are very popular as the shock-capturing methods of choice for the pure fluid Euler equations. In the first part of this talk I will discuss issues related to the extension of Godunov methods to solve the equations of compressible ideal MHD, and more in general issues related to the finite volume discretization of these equations (e.g., the $\text{Div}(\mathbf{B})$ problem). In the second part, I will present results from simulations performed with a higher-order Godunov method for the MHD equations. In particular, I will discuss the effects of these numerical methods on the physics of the system, and their possible interpretation as a subgrid-model for the MHD equations.

March 3, 1999

Preconditioning Constrained Systems

Andy Wathen

wathen@comlab.ox.ac.uk

Oxford University

Abstract

The general importance of preconditioning in combination with an appropriate iterative technique for solving large scale linear(ised) systems is widely appreciated. For definite problems (where the eigenvalues lie in a half-plane) there are a number of preconditioning techniques with a range of applicability, though there remain many difficult problems. For indefinite systems (where there are eigenvalues in both half-planes), techniques are generally not so well developed.

Constraints arise in many physical and mathematical problems and invariably give rise to indefinite linear(ised) systems: The incompressible Navier–Stokes equations describe conservation of momentum in the presence of viscous dissipation subject to the constraint of conservation of mass, for transmission problems the solution on an interior domain is often solved subject to a boundary integral which imposes the exterior field, in optimisation the appearance of constraints is ubiquitous.

We will describe two approaches to preconditioning such constrained systems and will present analysis and numerical results for each. In particular, we will describe the applicability of these techniques to approximations of incompressible Navier–Stokes problems using mixed finite elements and the MAC finite difference scheme. We will also present the results of computations on various matrices from the CUTE optimisation test set.

February 25, 1999

Large-Scale Parallel Unstructured Multigrid Computations for Steady-State Aerodynamic Problems

Dimitri Mavriplis

dimitri@icase.edu

NASA Langley Research Center

Abstract

Recent experiences in the development and execution of a parallel unstructured multigrid solver are discussed. The target applications consist of steady-state viscous turbulent flow over external aircraft configurations. The discretization is based on the vertices of an unstructured mesh that may contain arbitrary combinations of tetrahedra, prisms, pyramids and hexahedra. An FAS agglomeration multigrid algorithm has been developed for accelerating the solution to steady-state. This approach uses a graph algorithm to construct the coarse multigrid levels from the given fine grid, similar to an algebraic multigrid approach, but operates directly on the non-linear system using the FAS approach. Two preconditioning techniques are employed to relieve stiffness due to high-aspect ratio grid cells in the boundary layer regions, and the stiffness associated with regions of nearly incompressible flow.

The preconditioned multigrid solver is parallelized using the domain-decomposition approach with the MPI message passing interface. Good scalability for

medium size problems is demonstrated on the SGI Origin 2000 using up to 128 processors, and the Cray T3E using up to 512 processors. A large scale case using 25 million grid points (150 million degrees of freedom) is solved on a dedicated Cray T3E-1200e provided by Cray Research Inc. using up to 1450 processors. Near perfect scalability is demonstrated in going from 256 to 1450 processors on this machine, with the 1450-processor case producing a solution in about one hour of wall clock time in 500 multigrid cycles.

Future work will concentrate on parallelizing the pre-processing operations such as coarse multigrid level construction and grid partitioning which currently are performed sequentially on an SGI ORIGIN 2000. A case involving 100 million grid points (600 million degrees of freedom) should be feasible from a memory and cpu-time perspective on the large T3E configuration, or other ASCI machines, once the pre-processing bottlenecks are addressed.

February 23, 1999

Overlapping Grids Applied to Naval Hydromechanics

Anders Petersson

andersp@na.chalmers.se

Chalmers University of Technology

Abstract

The three-dimensional overlapping grid generator “Chalmesh” is presented. The code implements a combination of previous approaches implemented in, for example, PEGSUS and CMPGRD together with some new ideas.

The algorithm starts by forming a boundary description based on the overlapping surface patches. The union of the overlapping surface patches is used to cut the holes in the component grids by using a mark-and-fill technique. The grid cells intersected by the boundary are identified by an octree-based search technique and the grid points in those cells are marked as being inside or outside the boundary by employing a modified ray method. The holes are then filled starting from the outside points in the marked cells.

After the holes have been cut, the remainder of the grid points are classified by the CMPGRD approach as either interpolation, discretization, or hole points. Thereafter, the grid is trimmed to remove all unnecessary interpolation points to minimise the overlap between the components. Finally, the interpolation data in viscous boundary layer grids is compensated for boundary mismatch.

Examples will be given for two marine applications: a ship and a propeller, where these grids have been used to solve the Reynolds averaged Navier–Stokes equations.

February 19, 1999

The Role of Nonlinear Wavelet Approximation in Image Processing and Partial Differential Equation Solvers

Ronald DeVore
devore@math.sc.edu

University of South Carolina

Abstract

Wavelets and multiresolution give efficient decompositions of functions into simple building blocks. These decompositions can be utilized in the design of numerical algorithms for image processing or solving partial differential equations (PDEs). Much emphasis has recently been placed on nonlinear algorithms whose goal is to use adaptively selected terms of the wavelet decomposition to minimize computational cost. We shall discuss optimal methods for adaptively choosing the terms in a wavelet decomposition and then discuss some of their implications in image processing and PDE solvers.

February 10, 1999

Automatic Differentiation: A Tool for Computational Science

Paul Hovland

hovland@mcs.anl.gov

Argonne National Laboratory

Abstract

Derivatives play an important role in computational science. They are used in optimization, sensitivity analysis, the solution of nonlinear equations, and a variety of inverse problems.

Automatic differentiation (AD) provides a mechanism for transforming code for computing a function into code for computing that function and its derivatives. Unlike finite differences, the derivatives computed by AD are analytic, and do not suffer from truncation error. AD also avoids the time-consuming and error-prone task of coding derivatives by hand. We discuss how AD works, the tools available for its application, and techniques that enable a knowledgeable programmer to achieve improved performance.

February 9, 1999

A Negative Norm Least-Squares Approach to Div-Curl Systems in Two and Three Dimensions

James Bramble

bramble@math.tamu.edu

Texas A&M University

Abstract

In this talk, we will present a direct finite element approximation of div-curl systems. The discretization is based on a reformulation of the div-curl systems using a negative norm least-squares approach. The method is optimal with respect to both the order of the approximation of the finite element space as well as the regularity of the solution. The corresponding algebraic system is symmetric positive definite and, moreover, it is well conditioned and hence can be solved effectively by the conjugate gradient method, for example. The computation of the discrete negative norms only involves the preconditioners of second order finite element problems. In particular a multigrid or multilevel preconditioner can be used.

February 8, 1999

SailFlow: Computing Air Flow Around Yacht Sails with Overture

Cheryl Fillekes

sail@iconz.co.nz

Doyle Bouzaid Sailmakers, New Zealand

Abstract

SailFlow is a program to optimize sails for racing yachts in a commercial sail design environment now in its sixth month of development. The Overture framework was selected as the basis for the gridding and aeroelastic computations. Overture's object model allows the coherent and consistent derivation of grids specific to yacht sails and rigging from more general utility grids. Overture's overlapping grid technology allows multiple sails to be trimmed and re-trimmed relative to each other without having to re-grid the entire computational domain. Overture's support for new partial differential equations solver implementation allows SailFlow the flexibility of doing elastic computations on the same grids as the flow solvers are implemented on. This also allows the implementation of new turbulence models, such as the SST $k-\omega$ model. Finally, because Overture is based on A++/P++, a clear path for migration to parallel platforms is pre-established.

February 1, 1999

Viscous Fingering in Non-Newtonian Fluids

Petri Fast

fast@cims.nyu.edu

New York University

Abstract

Thin gap flows of non-Newtonian liquids are important technologically, and can exhibit a rich variety of dynamical behaviors. To study such problems, I use the thin gap limit to reduce a continuum viscoelastic model to a two-dimensional nonlinear generalization of Darcy's Law. This leads to a nonlinear elliptic boundary value problem to be solved in a time-dependent domain. Recent intermediate time-scale simulations of the resulting dynamical system show intriguing consistency with experiments. I present the current state of a project to develop a moving overset grid scheme for the accurate and efficient long-time simulation of this problem.

January 28, 1999

Analytical Performance Modeling of Hierarchical Mass Storage Systems

Yelena Yesha

yeyesha@cs.umbc.edu

University of Maryland and NASA Goddard

Abstract

Mass storage systems are finding greater use in scientific computing research environments for retrieving and archiving the large volumes of data generated and manipulated by scientific computations. This talk presents a queueing network model that can be used to carry out capacity planning studies of hierarchical mass storage systems. Measurements taken on a Unitree mass storage system and a detailed workload characterization provided by the workload intensity and resource demand parameters for the various types of read and write requests. The performance model developed here is based on approximations to multiclass Mean Value Analysis of queueing networks. The approximations were validated through the use of discrete event simulation and the complete model was validated and calibrated through measurements. The resulting model was used to analyze three different scenarios: effect of workload intensity increase, use of file compression at the server and client, and use of file abstractions.

January 25, 1999

Pattern Formation in Non-Newtonian Hele–Shaw Flow

Lou Kondic

kondic@math@duke.edu

Duke University

Abstract:

We explore the morphology of patterns due to the Saffman–Taylor instability in Hele–Shaw cells and find that it can be dramatically altered by the non-Newtonian response of complex fluids such as liquid crystals and polymer solutions. The dense-branching morphology of Newtonian liquids may be replaced by dendritic fingers with stable tips and sidebranches.

Starting from a very general viscoelastic fluid model, we find a distinguished limit where shear thinning effect is dominant. Darcy's Law leads to the nonlinear boundary value problem for the pressure in the fluid. Full numerical simulations show that shear thinning alone modifies considerably the pattern formation and can produce fingers whose tip-splitting is suppressed, in agreement with experimental results. These fingers grow in an oscillating fashion, shedding “side-branches” from their tips, closely resembling solidification patterns (*Phys. Rev. Lett.*, **80** 1998: 1433).

January 14, 1999

Domain Decomposition Techniques for Localized Space and Time Grid Refinement for Flow Problems

Richard Ewing
ewing@isc.tamu.edu

Texas A&M University

Abstract

We discuss various strategies for solving steady-state and transient problems with localized solution in space and time. The discretization techniques include Galerkin finite element and finite volume methods in space and discontinuous Galerkin and finite difference methods in time on locally refined (matching and non-matching) grids in both space and time. Next, we shall present iterative methods based on domain decomposition technique for solving the composite grid system. These include overlapping and nonoverlapping Schwarz methods (BPS and BEPS iterative algorithms) and domain decomposition methods for mortar approximations. Finally, we discuss numerical results demonstrating the capabilities of the methods for both test problems and applied problems from flows in porous media.

January 12, 1999

Domain Decomposition Preconditioners for High-Order Discretization Methods

Calvin Ribbens

ribbens@vt.edu

Virginia Institute of Technology

Abstract

A widely-used family of techniques for numerically solving elliptic partial differential equations (PDEs) on parallel processors is known as “domain decomposition” (DD). The key step in a DD method is typically the construction and application of a preconditioner for a Krylov subspace iterative linear equation solver. Most DD preconditioners are derived and analyzed in the context of second-order accurate discretization schemes, e.g., “five-point” finite differences, piecewise linear finite elements. In this talk, we describe two preconditioners for linear systems arising from higher-order discretization methods. We first describe HODIEX, a DD-preconditioned iterative solver based on the fourth-order accurate HODIE method (High-Order Difference Approximation with Identity Expansion) of Lynch and Rice. We then focus on ColEVP, a preconditioner for linear systems arising from a piecewise Hermite bicubic collocation discretization method (which also gives $O(h^4)$ accuracy).

The ColEVP preconditioner is defined in terms of a three-level grid and discretization scheme. In the framework of substructuring, we partition the domain into subdomains, edges, and vertices. The three-level discretization scheme has a fine grid G_h, h with uniform step h on the first level. The third level is the coarse grid G_H, H with step H corresponding to a subdomain width. The middle level is the edge-grid G_h, H' , a hybrid fine/coarse grid which has step h along each edge and step H' in the perpendicular direction. Here, $H' = H - h$ for all edges except those that are a distance H from a physical boundary—for those edges we use $H' = H$. The edge and vertex subproblems that must be solved in applying the preconditioner are based on G_h, H' and G_H, H , respectively. We show how the coupling between the edge and subdomain subproblems is significantly reduced by using the edge-grid G_h, H' , and how this can be used to construct an efficient preconditioner. We describe the motivation, construction, and implementation of ColEVP, and illustrate its performance with examples.

January 11, 1999

A Multigrid Approach for the Mortar Finite Element Method

Joseph Pasciak
pasciak@math.tamu.edu

Texas A&M University

Abstract

A multigrid technique for uniformly preconditioning linear systems arising from a mortar finite element discretization of second-order elliptic boundary value problems is described. These problems are posed on domains partitioned into subdomains, each of which is independently triangulated in a multilevel fashion. The multilevel mortar finite element spaces based on such triangulations (which need not align across subdomain interfaces) are in general not nested. Suitable grid transfer operators and smoothers will be discussed which lead to a variable V-cycle preconditioner resulting in uniformly preconditioned algebraic systems. Computational results illustrating the theory are also presented.

January 8, 1999

Spatial Databases and Geographic Information Systems

Hanan Samet

hjs@umiacs.umd.edu

University of Maryland

Abstract

An introduction is given to the spatial database issues involved in the design of geographic information systems (GIS) from the perspective of a computer scientist. Some of the topics to be discussed include the nature of a GIS and the functionalities that are desired in such systems. Representation issues will also be reviewed. The emphasis will be on indexing methods, as well as the integration of spatial and nonspatial data. An example spatial browser for a geographic information system using these concepts will be discussed.

December 18, 1998

Multi-Source Data Analysis in Science and Engineering

Samuel Uselton

uselton@nas.nasa.gov

MRJ Technology Solutions

Abstract

As digital data acquisition becomes easier, cheaper and more pervasive, and computational simulations gain increasing fidelity and detail, many activities can benefit from the combined analysis of data from several sources. Building a useful multi-source analysis system requires solving many problems, some of a pragmatic, engineering nature, and some of a more basic nature. The issues seem remarkably similar whether the application is weather modeling, environmental assessment and remediation planning, oil exploration and production, or engineering design processes. The problems to be addressed range from efficient access to large amounts of data from multiple heterogeneous sources, to design of user interfaces and inventing visualization techniques. Integrated display and comparative analysis of relevant data is interesting and relatively unexplored.

Our Multi-Source Visualization (MSV) project uses a specific problem domain, concurrent design of aircraft, to focus research and development efforts in this area.

Frequent contacts with the Earth Observation System (EOS) project, meso-scale atmospheric modeling researchers, and nanotechnology researchers, among others, keep broader needs in view.

All the work done by the Data Analysis Group of the NAS Division at NASA Ames Research Center is relevant and useful to this project. This group's work ranges from specific visualization techniques, through innovative user interfaces, the software engineering required to build complete systems, systems level performance improvements, effective access to very large data sets, and exploitation of large heterogeneous collections of scientific and engineering data. Much of this work will be touched upon, but the focus will be on the work that is specifically driven by the need for a variety of users to exploit common collections of data from many diverse sources.

This work is very much in progress. There are results to show, but also problems not yet satisfactorily resolved.

December 15, 1998

Does the Large Scale Determine the Small Scale in Turbulent Flows?

Heinz-Otto Kreiss

kreiss@math.ucla.edu

University of California, Los Angeles

Abstract

In meteorological applications, one can only measure the large scale. It is believed, however, that the small scale (mesoscale) plays an important role for longer-time forecasts. Since the observational net does not resolve the small scale, we want to discuss, for a model problem (2D and 3D Navier–Stokes equations), whether we can reproduce the small scale by the time history of the large scale.

December 11, 1998

New Algorithms for Incomplete Factorization: Symmetric Reduction Redux

Alex Pothen

pothen@cs.odu.edu

Old Dominion University

Abstract

We describe new algorithms for computing incomplete factor preconditioners for solving systems of linear equations by iterative methods. These algorithms rely on a newly developed structure theory for identifying “fill” in incomplete factorization. We use a technique called symmetric reduction to reduce the time needed to compute the preconditioners when high levels of fill are permitted. The new algorithms are provably faster than current implementations, and are also inherently parallel. Our implementations show that the new algorithms are faster than ILUT preconditioners that rely on a numerical threshold to compute the incomplete factors.

December 10, 1998

Some New Parallel Adaptive Finite Element Methods

Michael J. Holst
mholst@math.ucsd.edu

University of California, San Diego

Abstract

We discuss a new approach to using parallel computers with adaptive multilevel finite element methods. We propose an algorithm, for general elliptic equations on arbitrary polyhedral domains, which has the following interesting features.

- The algorithm begins with one processor solving a special small “coarse” problem, and then broadcasting the coarse problem and its solution to the remaining processors.
- The computation then proceeds independently on each processor, without communication (i.e., no boundary exchange communications are required).
- The final adapted mesh on each processor has nearly the same number of elements, so that load balancing is not necessary.
- The resulting solution is (provably) “as good” as the solution produced by the usual communication-intensive boundary exchange methods coupled with dynamic load-balancing strategies.

The parallel algorithm will be described in detail. To illustrate that the parallel algorithm can be used in conjunction with any sequential adaptive finite element method, we will give some examples using two different adaptive codes: the well-known 2D package PLTMG, and a new similar package called MC. MC is a dimension-independent, simplex-based, ANSI-C finite element code for covariant problems on d -manifolds ($d=2,3,\dots$). MC implements many features of the PLTMG, including adaptive error control, unstructured algebraic multilevel methods, global Newton methods, and continuation.

To illustrate the generality of the approach, the examples will include scalar linear elliptic equations in 2D, nonlinear elasticity in 3D, and the Hamiltonian and momentum constraints in the Einstein equations (a coupled four-component nonlinear elliptic system in 3D). We finish the talk by outlining the some local a priori and a posteriori error estimates of Xu and Zhou, which provide a theoretical justification for the new algorithm.

November 24, 1998

Advising the President on Information Technology

David M. Cooper
cooper31@llnl.gov

Lawrence Livermore National Laboratory

Abstract

This presentation focuses on the recommendations of the President's Information Technology Advisory Committee in the Interim report to the president submitted in August 1998. The major findings of the committee are that the federal investment in Information Technology (IT) R&D is inadequate and that the federal IT investment is too heavily focused on near-term problems. To offset these trends, the committee recommends that the United States government fund a strategic initiative in fundamental IT research with specific investments in software, high-end computing, scalable information infrastructure, and socio-economic and workforce activities. Specific details of each of these elements are presented.

November 6, 1998

Applications of ART Neural Networks

Gail Carpenter

gail@cns.bu.edu

Boston University

Abstract

ART (Adaptive Resonance Theory) neural networks for fast, stable learning and prediction have been applied in a variety of areas. Applications include airplane design and manufacturing, automatic target recognition, financial forecasting, machine tool monitoring, digital circuit design, chemical analysis, and robot vision. Supervised ART architectures, called ARTMAP systems, feature internal control mechanisms that create stable recognition categories of optimal size by maximizing code compression while minimizing predictive error in an on-line setting. Special-purpose requirements of various application domains have led to a number of ARTMAP variants, including fuzzy ARTMAP, ART-EMAP, ARTMAP-IC, Gaussian ARTMAP, and distributed ARTMAP. This talk will discuss some ARTMAP applications, including the following.

Multi-sensor fusion, with application to sonar target recognition

A new ARTMAP variant, called ARTMAP-FTR (fusion target recognition), has been developed for the problem of multi-ping sonar target classification (Carpenter and Streilein, 1998). The development data set, which lists sonar returns from underwater objects, was provided by the Naval Surface Warfare Center (NSWC) Coastal Systems Station (CSS). The ARTMAP-FTR network has proven to be an effective tool for classifying objects from sonar returns. The system also provides a procedure for solving more general sensor fusion problems.

Geospatial mapping from satellite remote-sensing data

A remote sensing testbed allows performance comparisons between neural network systems and state-of-the-art image processing and recognition techniques, in a collaborative project between researchers at the Center for Adaptive Systems and the Boston University Center for Remote Sensing (Carpenter, Gopal, Martens, and Woodcock, 1997). Network hierarchies that take advantage of the database size and structure have been developed. These systems provide an efficient method for producing accurate geospatial maps from high-dimensional satellite and terrain data. A new ARTMAP network is being considered as the candidate algorithm for the Moderate Resolution Imaging Spectrometer (MODIS) land cover product of the NASA Earth Observing System (EOS).

Self-organizing expert systems and computer-assisted medical diagnosis

Medical databases present many of the challenges found in general information management settings where speed, efficiency, ease of use, and accuracy are at a premium. A direct goal of improved computer-assisted medicine is to help deliver quality emergency care in situations that may be less than ideal. Working with these problems has stimulated a number of ART architecture developments in recent years, including ARTMAP-IC (Carpenter and Markuzon, 1998). A new collaborative effort brings together medical statisticians and clinicians at the New England Medical Center with researchers developing expert system and neural network learning systems.

November 5, 1998

Data Mining in Very Large Dimensional Data Sets

George Karypis

karypis@cs.umn.edu

University of Minnesota

Abstract

Data sets with high dimensionality pose major challenges for conventional data mining algorithms. For example, traditional clustering algorithms such as K-means or AutoClass fail to produce good clusters in large dimensional data sets even when they are used along with well known dimensionality reduction techniques such as Principal Component Analysis. Similarly, traditional classification algorithms such as C4.5 perform poorly on large dimensional data sets.

This talk presents a novel method for clustering related data items in large high-dimensional data sets. Relations among data items are captured using a graph or a hypergraph, and an efficient multi-level graph partitioning algorithm is used to find clusters of highly related items. We present results of experiments on several data sets

including S&P500 stock data for the period of 1994-1996, protein coding data, and document data sets from a variety of domains. These experiments demonstrate that our approach is applicable and effective in a wide range of domains, and outperforms techniques such as K-Means, even when they are used in conjunction with dimensionality reduction methods such as principal component analysis or latent semantic indexing scheme.

This talk also presents a graph-based nearest-neighbor classification scheme in which the importance of discriminating variables is learned using mutual information and weight adjustment techniques. Empirical evaluations on many real world documents demonstrate that this scheme outperforms state of the art classification algorithms such as C4.5, Ripper, Naive-Bayesian, and PEBLS.

November 4, 1998

Preconditioning Operators for Elliptic Problems with Bad Parameters

Sergei Nepomnyaschikh

Russian Academy of Sciences

Abstract

The construction of preconditioning operators for the iterative solution of systems of grid equations approximating elliptic boundary value problems with “bad” parameters is considered. These bad parameters characterize jumps in the coefficients, shape of the subdomains with composite materials, etc.

The suggested technique is based on the domain decomposition, i.e., the original problem is decomposed into subproblems in which the coefficients of equations inessentially change. We consider two variants of the domain decomposition technique: splitting into nonoverlapping subdomains and splitting into overlapping subdomains “without overlapping in the coefficients.” For the first domain decomposition method constructing of preconditioning operators involves constructing easy invertible norms for grid functions on the boundaries of the subdomains and norm-preserving explicit extension operators of func-

tions from the boundaries into inside subdomains. These norms should be equivalent to the norms which are generated by original problems on these boundaries (Schur Complements).

The second domain decomposition method is based on the theorem of extension of functions from subdomains with “strong” coefficients into subdomains with “weak” coefficients. Multilevel preconditioning operators for the subproblems on unstructured grids with the smooth coefficients are suggested and used in both methods. Design of these preconditioning operators is based on the fictitious space method. This method reduces the problem from the unstructured grid to a structured grid (at the first step) and then (at the second step) reduces the problem to a structured hierarchical grid where BPX-like preconditioners can be used. The convergence rate of the preconditioned iterative process is independent of the grid size and the bad parameters.

October 29, 1998

Parallel Particle-in-Cell Modeling of Semiclassical Quantum Models

Roy Hemker

University of California, Los Angeles

Abstract

We are modeling many-particle quantum systems by combining a semiclassical approximation of Feynman path integrals with parallel computing techniques previously developed at UCLA for simulating plasmas.

October 22, 1998

Clustering of Large Data Sets: The Possibilities Are Endless!

Wray Buntine

wray@ultimode.com

Ultimode Systems, LLC

and University of California, Berkeley

Abstract

Clustering or unsupervised learning is a basic form of pattern recognition developed to a fairly advanced state in the Sixties. Techniques for automatically selecting the “right” number of clusters, determining irrelevant features, and clustering of more sophisticated models have been in successful use for many years (despite claims by some that these are open research questions). Following well-publicised efforts at JPL some years back, the data-mining community is currently seeing a flurry of activity in scaling up clustering to large data sets.

First, I review potential methods available in the scientific literature, many of which are developed outside of the data-mining community and not well explored. Included are bump-tree and k-d tree techniques used for speeding up nearest neighbor algorithms, “clustering” methods from optimization in CAD, numerical analysis and high-performance computing (the term clustering here is used differ-

ently), automated software engineering methods for tuning on parallel processes and code-optimization, iterative, “hard” and deterministic-annealed versions of the underlying algorithms, and hierarchical clustering. Combinatorially, there are an enormous range of possibilities here and data mining researchers have barely scratched the surface. Thus I view the scaling up of clustering as largely a sophisticated software and engineering issue.

Second, I discuss some of our practical experiences in applying clustering for both commercial and government clients: We firmly believe intelligent preprocessing and post-visualization needs to be done by an innovative professional, and that the core software used in clustering sometimes need adaptation to individual problems, simple clustering models used for instance in SAS and standard earth science packages are too inflexible and the software itself may need to be customized.

October 20, 1998

Computational Fluid Dynamic Studies of Arterial Flow Disturbance Induced by Intravascular Stents

E. Tina Cheng
etcheng@ucdavis.edu

University of California, Davis

Abstract

Atherosclerosis is an arterial disease whose pathological complications, namely heart disease and stroke, are the leading causes of mortality in the industrialized world. In its advanced form, atherosclerosis leads to plaques which protrude into arterial lumens and form stenoses, or even complete vessel occlusions that obstruct blood flow and give rise to the pathological events. One common interventional procedure involves the placement of an intravascular stent, an expandable wire mesh structure that is introduced into the diseased artery in a compressed state and inflated at the stenosis or occlusion site to both restore blood flow and provide structural stability to the arterial wall. The major limitation to the success of this procedure, however, is restenosis, a complex and incompletely understood process by which plaques reform and re-protrude into the vessel lumen within a period of a few months.

The placement of a stent in an artery mechanically damages the endothelium, the monolayer of cells lining the inner surfaces of all blood vessels. In vitro data indicate that the rate of endothelial repair after injury may be significantly slower in regions in which endothelial cells are exposed to relatively large fluid mechanical shear stress gradients, as occurs at the end points of flow separation

zones. Therefore, flow separation in the vicinity of a stent may contribute to restenosis. Furthermore, clinical data suggest that the incidence of restenosis is higher for thicker-wire stents. We hypothesized that the occurrence of flow separation depends on hemodynamic matching between the stent wire thickness and the flow and geometric properties of the arterial segment in which the stent is positioned. To test this hypothesis, we have been studying the impact of stent wire thickness on the occurrence of flow separation under various arterial geometric and flow parameters using computational fluid dynamic techniques.

Our results to date have demonstrated that for straight arterial segments in steady flow, flow separation is more likely to occur in the case of thicker-wire stents and for higher flow Reynolds numbers. Flow pulsatility leads to periodic appearance and disappearance of flow recirculation zones thereby introducing significant temporal (in addition to the spatial) shear stress gradients. Finally, vessel curvature leads to additional flow recirculation zones not present in straight arterial segments. These results may provide insight into the mechanisms governing the clinical observation of increased incidence of restenosis in thicker-wire stents and may guide strategies for targeting particular stent structures for specific vascular sites.

October 19, 1998

Wavefunction Engineering of Quantum Devices Computational Issues in Modeling Nanostructures

Ramdas Ram-Mohan

lrram@wpi.edu

Worcester Polytechnic Institute

Abstract

Quantum semiconductor heterostructures provide new opportunities to investigate fundamental quantum mechanical effects while holding forth the promise of new optoelectronic devices. The talk will focus on the theoretical and computational issues that have to be addressed in understanding the optical properties of such structures. The efficacy of the finite element method developed by the speaker for the calculation of energy levels in quantum semiconductors is demonstrated, and their application to structures such as quantum wells, superlattices, quantum wires, and checker-board superlattices, and their optical properties are presented. Computational challenges in this area of quantum applications are as demanding as, for example, in fluid mechanics.

The new developments in computation have led to the concept of band-gap engineering of semiconducting materials being replaced by the new paradigm of wavefunction engineering for quantum semiconductor structures. The computational problems faced in solving sparse banded matrix equations and in determining eigenvalues of large sparse complex banded matrices will be considered. The need for parallel computing, the development of faster algorithms for parallel computers for matrix analysis, and the visualization of wavefunction calculations are some of the computational issues and future directions in quantum device modeling. This talk includes elementary conceptual introductions to the topics.

October 16, 1998

Stochastic Programming: A Model for Decision-Making under Uncertainty

Roger Wets

rjbw@math.ucdavis.edu

University of California, Davis

Abstract

Almost all (important) decision problems involve some level of uncertainty, either about data measurements, the values to assign to parameters describing future evolution or even about the environment in which one has to operate. A few typical examples are: trajectory selection with data uncertainty about the initial conditions; management of energy resources with uncertainty about hydro-power supply and prices; and groundwater remediation with uncertainty about the (nonhomogeneous) soil composition.

The main objective is to provide a brief introduction to stochastic programming models and solution techniques. Because the choice of an optimal (or a good) decision must take into account a huge number of possible realizations of the uncertain (stochastic) parameters of the problem, the development of algorithmic procedures must take advantage of the available architectures that are now becoming available.

October 13, 1998

Parallel, High-Resolution Finite Element Analysis of Trabecular Bone Biomechanics

Glen Niebur

gln@biomech2.ME.Berkeley.EDU

University of California, Berkeley

Abstract

We will describe ongoing research at the Orthopaedic Biomechanics Laboratory at UC Berkeley on development of spatially high-resolution (10–20 micron) finite element models of trabecular bone, the spongy bone tissue found in the spine and at all articulating joints such as the hip and knee. Trabecular bone is most involved in diseases such as osteoporosis and in prosthesis fixation.

The overall long-term goal of this work is to develop computational modeling techniques that can be used to investigate the effects of aging and disease on the mechanical behavior of this material. Over the past three years with funding from LLNL, we have successfully developed and optimized the parallel finite element codes to solve this class of problems, with models having up to 5 million elements. We now describe our work on the numerical convergence behavior of this technique, and applications to prediction of lower level material problems of this cellular network, as well as simulation of mechanical testing techniques, and investigation of failure behavior and mechanisms.

October 12, 1998

Algorithms for Rapid IsoSurface Extraction

Charles Hansen

hansen@cs.utah.edu

University of Utah

Abstract

Exploratory scientific visualization is a valuable paradigm for understanding complex physical phenomena. When these phenomena have associated volumetric (3D) scalar fields, isosurface extraction is a critical tool. An isosurface is the set of points where the scalar field has a particular value, the “isovalue.” The position of an isosurface, as well as its relation to other neighboring isosurfaces, can yield clues to the underlying structure of the scalar field. The ability to interactively change the isovalue provides insight into structure and underlying physical meaning of the scalar field. This talk will focus on the work being done at Utah for rapid isosurface extraction algorithms that enable exploration of large datasets for both local and remote visualization. These algorithms provide fast localization of an isosurface within very large datasets, as well as giving scientists prompt feedback during exploratory sessions. This work was the Ph.D. Dissertation topic for Yarden Livnat, who defended at Utah in 1998.

October 6, 1998

Parallel Multigrid Solvers for Finite Element Problems in Solid Mechanics

Mark Adams

madams@CS.Berkeley.EDU

University of California, Berkeley

Abstract

Prometheus is our parallel linear equation solver for 3D finite element problems on unstructured meshes that uses a classical (geometric) multigrid method with an algebraic architecture—that is, the user need only provide the fine grid. This talk discusses our methods of automatically constructing good coarse grids for solid mechanics problems, as well as methods for optimizing parallel performance of parallel multigrid codes and PETSc, useful on machines with relatively slow communication (e.g., a large machine with a good network or a small machine with high latency).

Some general issues of multigrid solver behavior on problems with large jumps in material coefficients and high Poisson ratio (incompressible materials) are also discussed. We will show numerical results for linear and non-linear problems in elasticity and plasticity, of up to 9,600,000 degrees of freedom (and perhaps 16.5 M degrees of freedom, with any luck), on a Cray T3E and an IBM PowerPC cluster.

October 5, 1998

Aspects of Robust Scalable Multigrid: Interpolation, Smoothing & Coarsening

Justin Wan

wlwan@math.ucla.edu

University of California, Los Angeles

Abstract

We discuss three major integrated components of multigrid to design robust multigrid methods. First, we describe a new approach to constructing robust interpolation operators, which can handle problems with problematic (discontinuous, oscillatory) coefficients in a unified fashion for both regular or irregular grids. The basic idea derives from recent domain decomposition theory, and is based on defining coarse basis functions that are stable (minimizing the total energy) and have good approximation property (preserving constants).

Second, we discuss the use of sparse approximate inverses as parallel smoothers for multigrid. They have been recently demonstrated to be useful technique for preconditioning in the parallel environments. We prove further that they are effective for eliminating local high-

frequency errors; hence they can be used as smoothers for multigrid. Furthermore, a distinctive feature over relaxation smoothers is that we may improve their quality by adjusting the nonzero patterns of the sparse approximate inverses, which is particularly useful for anisotropic problems.

Third, we describe a new geometric technique to construct coarse grids/ subspaces for discontinuous coefficient problems. The interface preserving coarsening selects coarse grid points so that all the coarse grids are aligned with the interfaces for regular interface problems on structured grids, and that the interfaces are resolved as much as possible for irregular interface problems. As a result, multigrid with simple linear interpolation is sufficient to obtain fast convergence.

October 1, 1998

Neutron Transport Computation with Variational Nodal Methods

Elmer Lewis

e-lewis@nwu.edu

Northwestern University

Abstract

A unified variational formulation of the even-parity form of the Boltzmann equation for neutrons is presented. Spherical harmonics (P_n), simplified spherical harmonics (SP_n) and discrete ordinates (S_n) approximations in angle reduce the transport equation to a set of coupled second-order partial differential equations, where the number of equations increases with the order n . These equations are formulated variationally, and a hybrid finite element method is applied to obtain spatial discretization which imposes particle balance over a coarse nodal mesh.

Two classes of trial functions are employed within the nodes: high-order orthogonal polynomials, and piecewise bilinear functions over finite subelements. In both cases changes of variables result in nodal response matrix equations. These are solved using red-black or multicolored iterative algorithms. Implementations of the P_n and SP_n approximations in the Argonne National Laboratory code VARIANT are discussed. New results are presented for piecewise bilinear trial functions; these allow each response matrix to include highly heterogeneous regions. The suitability of the methods for parallel computation is examined.



Institute for Scientific Computing Research

Visiting and Collaborating Professors' Project Abstracts



Applying ATLAS Technology to ASCI-Related Applications

Jack Dongarra

University of Tennessee

Abstract

Today's microprocessors have peak execution rates of gigaflop/s. However, straightforward implementation in Fortran or C of computations based on simple loops rarely results in such high performance. To realize such peak rates of execution for even the simplest of operations has required tedious, hand-coded, programming efforts.

In general, the existing BLAS have proven to be very effective in assisting portable, efficient software for sequential, vector and shared-memory high-performance computers. Hand-optimized BLAS are expensive, however, tedious to produce for any particular architecture, and in general will only be created when there is a large enough market, which is not true for all platforms. The process of generating an optimized set of BLAS for a new architecture or a slightly different machine version can be a time-consuming process. The programmer must understand the architecture, how the memory hierarchy can be used to provide data in an optimum fashion, how the functional units and registers can be manipulated to generate the correct operands at the correct time, and how best to use the compiler optimization. Care must be taken to optimize the operations to account for many parameters such as blocking factors, loop unrolling depths, software pipelining strategies, loop ordering, register allocations, and instruction scheduling.

We have developed a general methodology for the generation of the efficient linear algebra kernels. In our approach, we have isolated the machine-specific features of the operation to several routines, all of which deal with performing an optimized on-chip, cache contained, (i.e., in Level 1 (L1) cache) matrix multiply. This section of code is automatically created by a code generator that uses timings to determine the correct blocking and loop-unrolling factors to perform an optimized on-chip multiply. The user may directly supply the code generator with as much detail as desired (e.g., the user may explicitly indicate the L1 cache size, the blocking factor(s) to try, etc). If such details are not provided, the generator will determine appropriate settings via timings. Our approach, called Automatically Tuned Linear Algebra Software (ATLAS), has been able to match or exceed the performance of the vendor-supplied version of matrix multiply in almost every case.

This package has been designed to encode and efficiently take advantage of hardware information. Currently, our automated optimization techniques are applied to generate highly efficient implementations of dense linear algebra kernels such as the BLAS, since these subprograms are essential computational building blocks of many scientific computing codes, as well as heavily used general purpose software libraries such as LAPACK or PETSc. Our research is well advanced, and has demonstrated the

Applying ATLAS Technology to ASCI-Related Applications (continued)

capability of automatically generating, for a wide range of computers, a number of highly efficient kernels achieving performance comparable to, and often better than, hand-optimized codes (often written in assembler) tailored specifically for a particular architecture. These encouraging results are in addition immediately and directly reusable and applicable to software designed and written for distributed-memory concurrent computers, since these optimized kernels are also heavily used locally by each processor in many scientific computing application areas.

Determining and characterizing the effectiveness of run-time optimization techniques is particularly relevant to modern sparse linear algebra software libraries that tend to hide from the user the internal storage format. For instance, a certain storage structure may dictate a particular algorithmic approach. Once a general scheme of access has been found based on given a storage structure, one promising idea involves analyzing the sparsity pattern of the matrix operand to find places where the effi-

cient static optimization techniques may be re-used. There are numerous ways in which this can be done, and it is almost certain that there will be no provably best way. In this case, it will also be necessary to search the space of available options during run-time.

We are experimenting with a variety of techniques for optimizing sparse matrix vector multiplication to take instruction sets, functional units, and memory hierarchies into account. Sparse matrix-vector multiplication is of course the inner loop in any iterative solver, even multigrid, as it includes all the interpolation, restriction, and smoothing operations. The structural properties of the application lead to sparse matrices that feature a sufficiently regular pattern, so that the automatic optimization techniques already integrated in ATLAS can be successfully re-used and applied to generate the appropriate basic sparse linear algebra kernels needed in many applications. Our plan for achieving the necessary and exceptionally high degree of portability and optimization leverages the experience of our team in developing ATLAS technology.

Developing a Tuned Version of ScaLAPACK's Linear Equation Solver

Jack Dongarra

University of Tennessee

Abstract

The LINPACK benchmark has been used as a yardstick in measuring the performance of the Top500 installed high-end computers. This benchmark was chosen because it is widely used and performance numbers are available for almost all relevant systems. The approach used in the LINPACK benchmark is to solve a dense system of linear equations. For the Top500, the benchmark allows the user to scale the size of the problem, and to optimize the software in order to achieve the best performance for a given machine. This performance does not reflect the overall performance of a given system, as no single number ever can. It does, however, reflect the performance of a dedicated system for solving a dense system of linear equations. Since the problem is very regular, the performance achieved is quite high, and the performance numbers give a good check of peak performance of a system.

By measuring the actual performance for different problem sizes n , a user can get not only the maximal achieved performance R_{\max} for the problem size N_{\max} but also

the problem size $N_{1/2}$ where half of the performance R_{\max} is achieved. These numbers together with the theoretical peak performance R_{peak} are the numbers given in the Top500. In an attempt to obtain uniformity across all computers in performance reporting, the algorithm used in solving the system of equations must confirm to the standard operation count for LU factorization with partial pivoting. In particular, the operation count for the algorithm must be $2/3n^3 + O(n^2)$ floating point operations.

Drawing upon our years of experience with ScaLAPACK and BLACS development, we have developed a version of the benchmark based on the hardware of the ASCI Blue-Pacific. Collaborating with Andrew Cleary at LLNL, we have achieved a performance of 2.144 teraflop/s. One major obstacle in achieving maximal performance with the present ASCI Blue setup is that it involves using multiple machines for one run. This problem has been shown to cause a loss of efficiency of approximately 20% on the current benchmark code on ASCI Blue.

Research on Parallel Adaptive Finite Element Methods

Michael J. Holst

University of California, San Diego

Abstract

We are developing and implementing parallel algorithms for the adaptive solution of systems of partial differential equations (PDEs) using the finite element method (FEM). PDEs lie at the heart of many problems in scientific computing, as many physical laws are most conveniently expressed through them. Many PDEs are derived from variational formulations of physical problems. These equations are so large that only the most efficient algorithms can be employed, and only those that are scalable to massively parallel computers have any chance of success. Moreover, since it is often necessary to model complicated geometries, many PDEs are discretized on general unstructured grids and the solver must perform efficiently without structured geometric grid information.

We have focused on several fundamental issues arising in the parallel adaptive solution of linear and nonlinear elliptic PDEs. Adaptive meshing algorithms are critical to the successful solution of many classes of PDEs. One challenging problem in this area is to incorporate such adaptive algorithms into a parallel-computing environment, since the final mesh (and hence the load balance) is usually not known *a priori*. Our UCSD group has developed a procedure based on globally defined grids, each of which is fully refined in an exclusive local region, which leads to efficient parallel adaptive methods, using minimal communication and substantial reuse of existing quality sequential software.

Numerical Methods for Partial Differential Equations in Large-Scale Scientific Computations

Raytcho D. Lazarov

Texas A&M University

Abstract

Construction, analysis and numerical testing of efficient discretization techniques for solving elliptic partial differential equations that allow for parallel implementation are the foci of our research. This ultimately imposes domain decomposition-type algorithms, in which each subdomain is uniquely assigned to a processor. The case when the meshes do not align on the interfaces between subdomains is considered. This situation occurs when either coarsening or refinement is done independently and in parallel on the subdomains, generating grids that do not match along the interfaces between the subdomains. Mortar finite element techniques to glue the solutions across the subdomain boundaries have been employed. We also study a least-squares stabilization technique for solving advection–diffusion problems and problems of linear elasticity using the “minus one” norm inner product.

Work on the *hypr*e Framework and Molecular Simulations on Massively Parallel Processors

Calvin Ribbens

Virginia Polytechnic Institute and State University

Abstract

Our work has two separate foci: scalable linear solvers and computational materials science.

We are assisting in the design, development, implementation, and testing of the *hypr*e framework and library. Our contribution shows how a variety of domain decomposition preconditioners fit into the framework. Also, the ability of the larger Equation Solver Interface (ESI) standard to incorporate these preconditioned solvers will be considered. We are also studying the design and performance of output coefficient access with regard to the CoefficientAccess class of the current *hypr*e design. The output method on objects of this class allows preconditioners to read coefficients of a matrix, e.g., as is needed by ILU. The tradeoffs between efficiency (“How quickly can I get the data and how well do I use the memory hierarchy?”) and expressiveness and interoperability (“How big is the set of preconditioners that can be easily expressed?”) need more study. We are also studying the design of the ESI and *hypr*e in terms of support for two-level, inner-outer, preconditioned solvers. We are investigating new algorithms of this type that combine domain decomposition approaches with multigrid in an effort to improve the applicability of multigrid without sacrificing scalability.

We have been collaborating for more than two years with Diana Farkas of the Materials Science and Engineering Department at Virginia Tech on work focused on molecular simulations on massively parallel machines. The PI is working with LLNL’s Patrice Turchi to incorporate the new real-space Electronic Structure (ES) approach to study the interplay between chemical

order and topological disorder in complex bulk amorphous alloys. Specifically, we are modifying the existing $O(N)$ Tight-Binding Molecular Dynamics (TBMD) codes, already implemented on the T3E at NERSC, to account for the d-electron behavior in materials. This extension has already been done in an $O(N^3)$ serial version of the codes. The new MD codes will be implemented on ASCI Blue-Pacific. We will then extend the TBMD codes to a full spd-atomic orbital basis set and to multi-species (the extension has already been done in an $O(N^3)$ serial version of the codes). All of this is done with a view toward parallelizing the ES and Monte Carlo (MC) parts of the scheme on the ASCI Blue machine.

Once parallel versions of the codes have been developed, the goal is to perform simulations on large systems (1,000-10,000 atoms) using an extended version of the $O(N)$ TB-MD parallel codes. The current codes require on the order of tens of hours (on a 32-node T3E) to perform a structural minimization of a simple system containing about 1,000 atoms. Several iterations of this minimization are needed to study different structures with a specific thermal history. To store the wavefunctions and atomic coordinates at intermediate steps (for later analysis), and also the final configuration, 50 MB of disk space and 50 GB of storage are required. A single configuration on a 32 PE system (with 1,000 atoms) is typically of the order of 1 MB. Note that the TB-MD codes have already been parallelized on a Cray environment. The other codes (ES and MC) require parallelization to make the entire loop efficient. The real-space formalism will lead to a full implementation of the codes on MPP machines.

Scalable Linear Solvers for Partial Differential Equations

Jinchao Xu

Pennsylvania State University

Abstract

Our work concerns scalable linear solvers with special emphasis on multigrid/multilevel methods. Multigrid methods are deemed to be one of the most powerful methods for solving large-scale algebraic systems arising from the discretization of partial differential equations. Perhaps partially because of their mathematical complexity and problem-dependent tuning and performance, multigrid methods are underutilized in ASCI applications. One particular area of fruitful collaboration is the development of the so-called algebraic multigrid methods, in which our experience ranges from theoretical analysis to code development.

Our particular expertise is efficient methods for convection-dominated convection–diffusion problems, including new monotone finite element schemes, ordering algorithms for effective Gauss–Seidel iteration, and efficient multigrid algorithms for unstructured grids.

Another area of collaboration is adaptive grid techniques in finite element discretizations. Whereas adaptive finite element grid refinement has mostly been done in an “unstructured” fashion, “locally structured” refinement technique appears to have great potential for efficient finite element implementation on high-performance computers.



Institute for Scientific Computing Research

University
Collaborative
Research Program
Subcontract
Abstracts



Structured Adaptive Mesh Refinement Running on Multi-Tier Computers

Scott Baden

University of California, San Diego

Abstract

We are investigating scalable parallel programming methodology and performance tradeoffs arising in structured adaptive mesh refinement running on multi-tier computers. Multi-tier computers, such as the ASCI Blue-Pacific machine, employ symmetric multiprocessor (SMP) nodes. Their performance tradeoffs are more of a challenge to master than previous-generation single-tier computers, which are based on single-processor nodes.

The program leverages previous work on the KeLP system, a C++ framework tailored to multi-tier architectures. We have experience with a "communication aware" load-balancing strategy that includes communication costs as part of the workload. We are refining the load-balancing strategy developed earlier, and exploring locality-enhancing and latency-hiding techniques on the ASCI Blue-Pacific platform. The investigation will deliver a computational testbed permitting LLNL scientists to explore portable scalable implementations of adaptive mesh applications running on a variety of platforms of interest to the Laboratory. The principal investigator will collaborate with John May of CASC to carry out performance analysis of the software infrastructure and to establish contact with potential Laboratory users.

Computational Fluid Dynamic Studies of Arterial Flow Disturbance Induced by Intravascular Stents

Abdul Barakat

University of California, Davis

Abstract

Atherosclerosis is an arterial disease whose pathological complications, namely heart disease and stroke, are the leading causes of mortality in the industrialized world. In its advanced form, atherosclerosis leads to plaques that protrude into arterial lumens and form stenoses, or even complete vessel occlusions that obstruct blood flow and give rise to the pathological events. One common interventional procedure involves the placement of an intravascular stent, an expandable wire mesh structure. The stent is introduced into the diseased artery in a compressed state and inflated at the stenosis or occlusion site to both restore blood flow and provide structural stability to the arterial wall. The major limitation to the success of this procedure, however, is restenosis, a complex and incompletely understood process by which plaques re-protrude into the vessel lumen within a period of a few months.

The placement of a stent in an artery mechanically damages the endothelium, the monolayer of cells lining the inner surface of all blood vessels. In vitro data indicate that the rate of endothelial repair after injury may be significantly slower in regions in which endothelial cells are exposed to relatively large fluid mechanical shear stress gradients, as occurs at the end points of flow separation zones. Therefore, flow separation in the vicinity of a stent may contribute to restenosis. The broad hypothesis driving our research is that the occurrence of flow separation depends on appropriate hemodynamic matching

between the stent design and the flow and geometric properties of the arterial segment in which the stent is positioned. We have been testing this hypothesis by using computational fluid dynamic techniques to probe the impact of various geometric and flow parameters on the occurrence of near-stent flow separation. Our results to date indicate that relatively small variations in stent wire thickness may have profound implications on the flow field within an artery and that thicker-wire stents are more likely to induce regions of flow separation with accompanying high wall shear stress gradients. This may partially explain the observation in animal studies of higher incidence of restenosis when using thicker-wire stents. We are extending the simulations to incorporate additional features that will render our numerical model more physiologically realistic. Our specific aims are (1) to perform three-dimensional steady and pulsatile flow simulations of the flow field in the vicinity of an intravascular stent positioned within a rigid-wall arterial segment with out-of-plane curvature, (2) to incorporate non-Newtonian fluid properties into the simulations, and (3) to incorporate certain aspects of arterial wall motion into the simulations.

The proposed collaboration with LLNL is critical for two reasons: LLNL's supercomputer are needed; and hydrodynamic simulation capabilities currently under development at LLNL may, in the future, allow the incorporation of additional physiological considerations into the aortic model, including fluid-wall coupling.

Implementing the Basic Self-Organizing Map (SOM) Algorithm for Data Mining

Jackson Beatty

University of California, Los Angeles

Abstract

We are developing a parallelized implementation of the basic self-organizing map (SOM) algorithm for data mining applications, designing a visually more effective SOM interface, and exploring the usefulness of volume SOM reference vector arrays. This research is important because SOMs have recently proven to be extraordinarily valuable tools for data mining in fields as diverse as mapping the distribution of atomic and molecular ions (Wolkenstein et al., 1997), identifying neuroactive compounds (Bucknecht et al, 1996), protein sequencing (Hanke et al., 1996) and galactic morphology (Naim et al., 1997).

We are implementing SOM procedures to study the problem of functional cortical organization by studying regularities in patterns of cerebral activation as evidenced in positron emission tomography and functional magnetic resonance imaging data. We believe that high-resolution self-organizing maps are best suited for our application, but high resolution can be produced only by using very large reference vector arrays. These are impractical using conventional serial computers, but are ideally suited for parallelization. We believe that parallelized SOM analysis, with improved visualization of the resulting SOM configurations, will not only materially aid us in our study of the functional cortical circuitry of the human brain, but will also contribute substantially to the computational infrastructure in a number of other scientific disciplines.

Parallel PIC Modeling of Semiclassical Quantum Models

Viktor Decyk and John M. Dawson

University of California, Los Angeles

Abstract

We are modeling many-particle quantum systems by combining a semiclassical approximation of Feynman path integrals with parallel computing techniques previously developed at UCLA for simulating plasmas.

Parallelized 3D Relativistic PIC Code for the Production of Useful Electron Bunches

Warren Mori

University of California, Los Angeles

Abstract

We are developing a parallelized three-dimensional relativistic particle-in-cell (PIC) code for studying the production of useful electron bunches using ultra-intense lasers-pulses. The electron bunches are produced when the radiation pressure of an intense laser either directly accelerates electrons or excites a plasma wave wake, which accelerates electrons. Ultra-short electron bunches could be useful for injection into high-energy particle accelerators, for radiation sources, and for the fast ignitor fusion concept. A reliable and robust code will be invaluable towards the realization of such ultra-short electron bunches in a laboratory.

We have carried out large physics simulations on the T3E at NERSC using an already existing code, PEGASUS, and have successfully developed a Fortran90-based object-oriented production code using the algorithms from PEGASUS. We will study the performance of a 3D parallelized version of this code and use it to investigate the production of useful electron bunches.

Algorithms and Software for Sensitivity Analysis of Large-Scale Differential–Algebraic Systems

Linda Petzold

University of California, Santa Barbara

Abstract

Sensitivity analysis of large-scale differential–algebraic (DAE) systems is important in many engineering and scientific applications. Sensitivity analysis generates essential information for design optimization, parameter estimation, optimal control, model reduction, management of uncertainty, process sensitivity and experimental design.

We are investigating and developing sensitivity analysis methods for large-scale DAEs. Our goal is the development of robust and efficient sensitivity software, based on the most recent versions of our widely employed codes DASSL and DASPK, and incorporating automatic differentiation software for evaluation of sensitivity derivatives. The proposed work is being done in close collaboration with Peter Brown and Alan Hindmarsh of LLNL, and includes both development and analysis of numerical methods and development of software.

Seismic Scattering for Strong Ground Motion Applications

Fred Pollitz

University of California, Davis

Abstract

Prediction of strong ground motion using forward computation of the propagating seismic wavefield is one of the few means by which to estimate the potential hazards posed by a future large earthquake. The most damaging ground motions typically have a period of 4 seconds (frequency 0.25 Hz). Methods to calculate seismic wavefields at such frequencies have continually evolved as the model used to represent Earth's crust and upper mantle has been refined. Current efforts to handle the complex scattering and wave propagation phenomena occurring between seismic source regions and Earth's surface rely primarily on finite difference or finite element computation. These methods are naturally able to account for general variations in elastic properties and are relatively simple to implement, but the amount of space demanded by them becomes prohibitive as the domain of wave propagation becomes larger.

We approach the problem of seismic scattering for strong ground motion applications using the analytic framework of coupled mode theory originally developed to perform very-low-frequency (0.01 – 0.05 Hz) seismic

wavefield computation. We have recently adapted these methods to higher frequency by introducing a refined technique for calculating the required wavefields, and are implementing them on LLNL supercomputers to simulate strong ground motion from crustal earthquakes, including the main effects of scattering through strongly heterogeneous crust and upper mantle structures. Our code can currently handle 3D variations in isotropic seismic wave velocities and density as well as attenuation. Previous research efforts in seismic scattering phenomena at lower frequency demonstrate the importance of including multifold mode-mode interactions in order to account for complex wave conversion effects over a broad range of spatial scales, as well as the importance of multiple scattering when large scattering domains are involved. Our work so far shows that these considerations are even more important at higher frequency. We intend to apply our new code to predict strong ground motion for various faulting scenarios in northern California, where combined body-wave and surface-wave based tomography now supply us with a satisfactory image of the highly variable crust and upper mantle structure.

Global Simulation of the Earth's Magnetosphere With Adaptive Mesh Refinement

Joachim Raeder

University of California, Los Angeles

Abstract

Our work adds structured adaptive mesh refinement capabilities to the UCLA Global Geospace Circulation Model (UCLA-GGCM) in collaboration with the SAMRAI project in the Center for Applied Scientific Computing at LLNL. The UCLA-GGCM is a global circulation model of Earth's magnetosphere and ionosphere that has been used for several years to study the interactions of the solar wind, magnetosphere, and ionosphere. The UCLA-GGCM with local adaptive mesh refinement supports a grid resolution in critical regions that is about two orders of magnitude better than the original code. This additional resolution allows us to address problems that were previously elusive to global modeling, such as flux transfer events, the formation of boundary layers, and the formation of thin current sheets in the late substorm growth phase.

SAMRAI is a general software support framework for structured adaptive mesh refinement (AMR) applications on parallel high performance computing hardware. To support local time refinement in the AMR time integration algorithm, we are investigating the numerical issues involved in local time synchronization of the magnetic and electric fluxes at coarse-fine grid interfaces.



Institute for Scientific Computing Research

Laboratory Directed Research and Development Project Abstracts



Novel Parallel Numerical Methods for Radiation and Neutron Transport

Peter N. Brown

Center for Applied Scientific Computing

Abstract

We propose to design advanced numerical methods for the parallel solution of three-dimensional radiation and neutron transport problems on massively parallel computers. Our emphasis will be on the development and implementation of methods for the parallel solution of the Boltzmann transport equation, as well as related equations.

Recent developments in the areas of First-Order System Least-Squares (FOSLS) methods show great potential for providing more accurate and robust solution procedures than current approaches. The FOSLS-based approach also provides a natural scalable multilevel (i.e., multigrid) solution procedure for the resulting discretized problems, and is independent of the discretization used in phase space. We also propose to develop more accurate phase space discretization techniques, in space and direction (angle) as well as better time-stepping algorithms.

Large-Scale Scientific Data Analysis and Visualization

Mark Duchaineau

Center for Applied Scientific Computing

Abstract

We propose a collaboration among CASC, University of California, Davis, University of Texas at Austin and LLNL Defense and Nuclear Technologies Department (DNT) scientists to develop and apply efficient multi-resolution and compression methods for high-performance visual exploration and precise quantitative analysis of extremely large scientific data (terabytes to petabytes).

In particular, we will investigate sparse (i.e., output-sensitive) wavelet transforms applied to variables defined over time-dependent, unstructured, adaptive 3D grids of mixed materials with general geometries, and fully output-sensitive, view-dependent optimization applied to the families of surfaces derived from such variables and grids. The algorithms will be fully scalable in the algorithmic sense, and efficiently implemented on distributed- or shared-memory parallel systems. Initially, the algorithms will be applied to ASCI and DNT physics simulations; a broad range of further applications is well known and will be aggressively pursued.

This project leverages considerable ASCI and Numerical Environment for Weapons Simulations (NEWS) funding, and is synergistic with other current and proposed CASC projects.

Applying Data Mining Technologies to Large-Scale Scientific Data Sets

Chandrika Kamath

Center for Applied Scientific Computing

Abstract

There is a rapidly widening gap between data collection capabilities and the abilities of scientists to explore, analyze, and understand this data. As a result, scientists at the Laboratory are expressing concern at the potential loss of useful information in this data. To address these problems, a new generation of computational techniques is needed to help automate the exploration and analysis of large-scale scientific data. Our research focuses on applying and extending ideas from the area of data mining, in particular pattern recognition, to improve the way in which scientists interact with large, multi-dimensional, time-varying data. As pattern recognition algorithms are common across problem domains, it will be possible to apply our research to many different applications across the Laboratory.

Data mining, though a field in its infancy, has shown promise of significant payoffs when applied to small, low-dimensional data. Several concerns, however, remain to be addressed in extending these algorithms to large-scale data, especially in the scientific domain. Our approach to scaling pattern recognition techniques to terabyte and petabyte data sets addresses the following issues.

- Effective preprocessing of the data using techniques such as dimension reduction and sampling to make the search for patterns in large, multi-dimensional data sets tractable.
- Development of new techniques to improve the effectiveness of pattern recognition algorithms.
- Error modeling to allow user control of accuracy versus time tradeoffs.
- Efficient parallel implementations to enable interactive exploration of data.

This research will combine advances in high-performance computing with techniques from machine learning, soft computing, and statistics. We will first demonstrate large-scale pattern recognition techniques in two test-bed applications, MACHO and FIRST. Then, in the follow-on years of this effort, we will apply our research to other relevant applications at the Laboratory, including ASCI.

Proposed work and anticipated results

During FY99 we propose to investigate several research ideas in image processing and pattern recognition algorithms. Using the image and catalog data from the FIRST project, we will

- Implement an algorithm to group radio sources.
- Create training and test data from existing known images of bent-doubles and sources other than bent doubles.
- Implement classification algorithms such as neural nets and decision trees to automatically identify bent doubles.
- Update the image catalog with additional data derived from the images.

We will also preprocess a subset of the MACHO images using image-processing techniques such as wavelets for de-noising and feature detection to generate an initial set of features that identify an asteroid.

For our future work, we will refine our image processing techniques as well as the clustering and classification algorithms using these initial applications, implement parallel versions in an object oriented framework, and couple our algorithms to the research on error modeling techniques.



Institute for Scientific Computing Research

Student Internship Research Summaries



Domain Decomposition Models for Parallel Monte Carlo Transport

Henry J. Alme

University of California, Davis

Research Summary

Monte Carlo methods are a desirable way to solve transport problems. Monte Carlo transport methods often provide the best answer for a particular problem, and—in some cases—are the only method available. The objection to Monte Carlo is the large computation time often required for the many samples that must be taken.

Parallel computers have great promise in increasing the performance of scientific computing, but applying parallelism to Monte Carlo transport is problematic. Monte Carlo has unique features that make it more difficult to parallelize when domain decomposition is required (by memory constraints, for example). One such feature is that the computational work for a Monte Carlo problem is localized; it is confined to the parts of the mesh where the sample particles actually travel. This makes conventional parallelization methods less useful.

This work proposes a strategy for parallelizing Monte Carlo transport problems that combines full replication (the most desirable strategy) with domain decomposition, (often required by memory constraints in state-of-the-art problems in scientific computation). The strategy uses various heuristics to estimate which areas of the computational mesh will require the most work, allowing those areas to be replicated more often than areas with less activity. This strategy was applied to several problems, some of them small test problems and some on the scale of real problems. While it was problem-dependent, the scaling behavior, the incremental performance improvement gained by adding more processors to a calculation, of the strategy presented in this work was generally superior to conventional domain decomposition, running two to three times faster in some cases. It required less communication and had better load-balance characteristics.

A Least-Squares Finite Element Method for 3D Neutron Transport Problems

Travis M. Austin

University of Colorado, Boulder

Research Summary

Application of a least-squares technique to the 3D neutron transport equation is an exciting alternative to current computational methods. Our method relies on a properly scaled least-squares functional to develop a minimization principle, which is lacking in a standard formulation. The scaling is designed to capture the correct behavior of the solution in different parameter regimes. The minimization principle provides a trusted framework for building the components of an optimal multilevel solution procedure.

Using a finite expansion of spherical harmonics to represent the angular variable leads to a self-adjoint, positive definite system of moment equations. A conjugate gradient method with a block Jacobi preconditioner, that employs multigrid to invert the diagonal elements, is used to approximately invert this system of equations. For nearly all parameter regimes and all diagonal elements, standard multigrid is sufficient as a preconditioner. For certain parameter regimes, however, multigrid performs poorly on the diagonal elements associated with the first-order moments. A new multigrid algorithm has been developed that simultaneously address the three first-order moments.

Our goal is to develop a multigrid routine that has convergence factors independent of discretization size. Smoothing works only on eigenvectors at the high end of the spectrum. Usually, these eigenvectors are exactly the oscillatory errors. With our operator, the other oscillatory

errors at low end of spectrum are left unattended. When the smoothing does not truly smooth, the coarse-grid correction becomes worthless as a component in the multigrid algorithm.

For the first part of my summer, I put the finishing touches on code that addressed this previously defined system of equations in 2D. The idea of the algorithm is to augment the standard multigrid algorithm with a special smoothing step. For a two-grid algorithm, we typically have a smoothing followed by a coarse-grid correction and then another smoothing. The new algorithm adds relaxation over divergence-free errors with each step of smoothing. This will “catch” the oscillatory components hiding at the low end of the spectrum. The algorithm we have developed results in convergence factors independent of discretization parameters and bounded by 0.6, which is very good reduction for such a complicated system.

The second half of my summer was spent building the components of the code to solve the full 3D neutron transport problem formulated with least-squares principles. Because of the complication of the new algorithm, particularly with the new finite elements we must introduce, issues of matrix storage, message passing, and efficiency of the algorithm arise. Having settled matrix storage and message passing issues, we have begun to develop the necessary components to extend my 2D algorithm to 3D as a preconditioner in the transport code, which is an ongoing project.

Symmetric Wavelets on Arbitrary Topology

Martin Bertram

University of California, Davis

Research Summary

Wavelets are used in lossy and lossless compression schemes for scientific data obtained from flow field simulations on supercomputers. Compression is necessary to store, transmit and visualize massive-volume datasets. Most compression schemes are restricted to data defined on regular, rectilinear grids. For compact representation of isosurfaces and material boundaries of arbitrary topology, however, wavelet compression schemes need to be generalized to handle data defined on polygonal meshes.

We constructed new wavelet basis functions defined on subdivision surfaces that have a polygonal base mesh and generate subgrids of regular topology, such as Catmull–Clark subdivision surfaces. Compared to most wavelet approaches on arbitrary topology, our wavelets generalize tensor product basis functions. As a conse-

quence, we obtain B-spline representations for reconstructed data, except in the neighborhood of a few extraordinary points in the base mesh.

We implemented two new wavelet transforms that generalize bilinear and bicubic B-splines to mesh domains of arbitrary topology. To obtain compression, we encode the wavelet coefficients that are sparse or have small absolute values, using a coding scheme that we developed earlier. The overall compression algorithm is highly efficient and provides surface reconstructions at multiple levels of accuracy.

We will use the new surface compression scheme for efficient extraction and visualization of isosurfaces. We also want to develop similar techniques for compression of volumetric data defined on irregular polyhedral meshes.

Material Interface Reconstruction

Kathleen S. Bonnell

University of California, Davis

Research Summary

Material interface reconstruction from volume fractions has been approached mainly through the discipline of computational fluid dynamics. We want a generalized visualization solution that can be used for time-varying data sets.

From a rectilinear data set containing m different materials, we want to reconstruct the boundary surface of each material. The problem is under-constrained: For general arrangements of volume fractions there is an infinite set of boundary surfaces that are consistent with the fractions. We know that each cell contains a numeric tuple that represents the volume fractions of each material in the cell, and whose sum is 1. Our goal is a simple and robust solution to visualizing the interface.

The general approach is to march through the data in a manner similar to marching cubes, calculating where the interface intersects cell edges, then triangulating these intersection points. Unlike the marching cubes approach, however, which relies on single scalar values for cell vertices, we have had to design an approach that will work with the volume fractions of our data. From our cell-centered representation of the data, we want to create a vertex-centered representation in order to utilize established methods of surface reconstruction where possible. By taking the center of each cell and assigning that cell's volume fractions to the vertex, we can work with the dual of the data set. We then subdivide each dual cell into 6 tetrahedra to reduce the number of cases to be considered for triangulation purposes. From here,

we march through each tet, calculating intersection points of the material interfaces. This is accomplished by mapping the volume fractions of the tet vertices into barycentric space. From barycentric space, we calculate material boundary intersections with edges and faces. Finally, we triangulate for each type of boundary: e.g., if we have three materials, a, b, and c, then we triangulate the interfaces between a and b, b and c, and c and a. We expect to get a continuous interface, with coherence over time (when data is coherent over time).

The original approach was proposed by Ken Joy for 2D. I spent the summer working on the 3D solution, searching for a simple enumeration scheme for the different cases (combinations of tet edge intersections and face intersections). We want to create a look-up table similar to marching cubes, one for which we know that no cases have been overlooked and one that can, we hope, be generalized for more than three materials. I have used small datasets to test the discovered cases, and the orientation of the triangulations. This portion is still a work in progress, although only a few more cases remain to be enumerated.

The ASCI visualization team will integrate our work into the terascale browser. I will continue to enumerate the triangulation cases for the three-material solution, solve the problem for more materials, and generalize the current method. I will then compare resulting volume fractions with original data for error analysis, and attempt to solve the problem for unstructured data and for an arbitrary number of materials.

Implementation of an Operator-Split Preconditioner

Thomas A. Brunner

University of Michigan

Research Summary

We would like to solve a system of differential–algebraic equations derived from a finite element discretization of the neutron transport equation. When solving this system, a preconditioned GMRES method is used. We need some way to approximate the iteration matrix as a preconditioner. The conventional preconditioner ignores the scattering. For some problems, the scattering can be very important and should be included somehow. The iteration matrix can be inverted in several stages, each stage dealing with a small part of the physics. This operator-split method was implemented in a time-dependent neutral particle transport code.

In order to test the new preconditioner, several simple test problems were devised. In all cases, a small cube with a side of length 1.0 was split up into four zones. Simple one-group isotropic scattering was tested first. The amount of absorption was slowly turned off. A parameter varying the diffusiveness of the problem was var-

ied. The low-diffusion case was difficult to solve, and both the original and the operator-split preconditioner had many linear convergence failures. A diffusion-based preconditioner was tried. Such a preconditioner often, but not always, outperforms the original. A five-group problem with down scattering was also tested. Here, the operator-split preconditioner performs well. Though it converged more rapidly, it also consumed more CPU time than the original.

The operator-split preconditioner is derived with the assumption of small time steps. In the limit where the time steps become vanishingly small, the scattering correction disappears and we recover the original preconditioner. It would be possible to use the original preconditioner for small time steps and to switch to the new one adaptive for longer steps. The limit of large time steps, however, is difficult for both preconditioners, meaning that further work is required for steady-state problems.

Coloring Algorithms for Element-Interpolation-Based Algebraic Multigrid

Tim Chartier

University of Colorado

Research Summary

Modern simulations often involve the solution of problems with billions of variables, mandating the use and development of scalable algorithms. While multigrid methods are scalable for many regular grid problems, they can be difficult to develop for the large unstructured grids that many simulations require. Algebraic Multigrid (AMG) attempts to overcome this difficulty by abstracting multigrid principles to an algebraic level so that the algorithm is more automatic and robust. Many researchers see AMG as one of the most promising methods for solving large-scale problems.

My research at Lawrence Livermore National Lab applied AMG to problems of elasticity with complex geometry. This work applies to energy research and the area of nuclear weapons stockpile stewardship. To protect our environment, we must understand how these materials behave in storage over long periods. Acquiring such knowledge is part of the mission of the Department of Energy's ambitious ASCI project.

Lawrence Livermore National Lab researchers are studying AMG in the context of the DOE ASCI project. My research plays a role in their efforts to create scalable algorithms for energy problems. In particular, I am researching coloring algorithms for AMG based on element interpolation (AMGe).

This summer, I continued research on exploiting features available in AMGe to select coarse grids to allow effective interpolation using element interpolation. I continued discussion with LLNL researchers who are also investigating and testing ideas in this area. I developed a very effective 2-level algorithm. Our current research focuses on developing a multilevel algorithm that maintains low element complexity, grid complexity, and convergence factors. It is difficult to balance these competing factors; however, current research shows promise.

Curvature-Based Adaptive Mesh Refinement for Level Set Tracking

Paul Covello

University of California, Davis

Research Summary

Level set tracking is used to compute burn tables to be used for detonation shock dynamics code. The level set represents the detonation shock front, which is assumed to propagate like an optical wavefront and thereby follows the eikonal equation. The advantage of using burn tables is that they greatly speed up and simplify the detonation shock dynamics code.

Algorithms exist to track level sets. When given a grid, an initial level set, and speed of detonation shock propagation these numerical codes can predict how the level set will propagate through the rest of the given domain. The advantage of these methods is that they are fast and they work on unstructured grids. The disadvantage is that when the curvature of the level set is high, the accuracy of these algorithms goes down. The way to remedy this is to have higher resolution, or more discrete points, where the curvature of the level set is high.

I am developing an algorithm that keeps a watch on the curvature of the level set: Where the curvature exceeds a certain threshold it will refine that portion of the grid. This new algorithm will be an add-on to the current level set tracking algorithms and also keep the advantage of computational speed and independence of grid type.

Thus far, I have achieved great improvement in accuracy via curvature-based adaptive mesh refinement for non-orthogonal structured grids in two dimensions. My next goal is complete unstructured grids in two dimensions and then do the same in three dimensions.

Large Eddy Simulation of Rayleigh–Taylor Instability Using the Arbitrary Lagrangian–Eulerian Method

Rebecca Mattson Darlington

University of California, Davis

Research Summary

This research addresses the application of a large eddy simulation (LES) to Arbitrary Lagrangian Eulerian (ALE) simulations of Rayleigh–Taylor instability. First, ALE simulations of simplified Rayleigh–Taylor instability are studied. The advantages of ALE over Eulerian simulations are shown. Next, the behavior of the LES is examined in a more complicated ALE simulation of Rayleigh–Taylor instability. The effects of eddy viscosity and stochastic backscatter are examined. The LES is also coupled with ALE to increase grid resolution in areas where it is needed. Finally, the methods studied above are applied to two sets of experimental simulations. In these simulations, ALE allows the mesh to follow expanding experimental targets, while LES can be used to mimic the effect of unresolved instability modes.

Classification of FIRST Data Using Support Vector Machines

Matt Giamporcaro

Boston University

Research Summary

With the proliferation of information-gathering and -processing resources, scientists have become flooded with more data than they can manage using traditional techniques. The Center for Applied Scientific Computing is developing new algorithms for the exploration and analysis of these large, complex data sets. By applying and extending ideas from data mining and pattern recognition, researchers on Project Sapphire are assembling a toolbox of computational techniques to help scientists extract useful information from these data.

The Support Vector Machine (SVM) is an algorithm that has recently received considerable attention in the pattern recognition literature. It has the potential to be a powerful tool for solving data classification and regression problems. This project has consisted of researching the theory behind SVMs and evaluating how they might be applied to the classification of large data sets. The test data set used has been the FIRST (Faint Images of the Radio Sky at Twenty-cm) set of radio-astronomical data.

The code for a SVM was written in C++. The specific training algorithm implemented is based on an approach presented by Platt (1998) called Sequential Minimal

Optimization. The SVM was applied to a subset of the FIRST data set and used to determine whether galactic radio sources might belong to an interesting class known as bent doubles.

The SVM was trained on the data, and test results were compared to those of a decision-tree algorithm implemented using an off-the-shelf software package (C5.0). In training, the performance of the SVM was similar to that of C5.0, in both run-time and cross-validation error rate. During testing on unknown data, there were significant differences in classification performance between the two systems. These results have not yet been analyzed to determine whether they represent better or worse performance on the part of the SVM.

Through the upcoming year, project researchers will continue to evaluate the feasibility of applying SVMs, as well as related techniques such as artificial neural networks, to this classification problem. Directions for continued research on SVMs include: determining whether the SVM performed better than C5.0; exploring the effect of parametric changes to the initial algorithm; and extending the published SVM algorithm to “less crisp” training data.

Explorations of a Stationary Diffusion Equation

Ana Iontcheva

University of California, Davis

Research Summary

I worked on three main problems during my LLNL research visit. First, I solved a stationary diffusion equation with the diffusion coefficient represented as a two-dimensional function and also as a matrix. I employed a preconditioned conjugate gradient method with the multigrid V-cycle as a preconditioner, and Gauss–Seidel as the pre-smoother and the post-smoother (forward and backward, respectively).

The second problem was to approach this same diffusion problem, using the classical hierarchical basis.

The third problem was to apply to the solution of this problem the Approximate Wavelet-Modified Hierarchical Basis presented in the paper “Stabilizing the Hierarchical Basis by Approximate Wavelets” by Panayot Vassilevski and Junping Wang. I finished the implementation (in C++) for these three methods, created visualizations, and am comparing their effectiveness.

A GUI Configure System for the KULL Code Packages

Mark Jeunnette

University of Chicago

Research Summary

The KULL code is a group of packages, written in FORTRAN, C, and C++, which together work to produce computational simulations of experiments to be carried out at the National Ignition Facility (NIF). Currently, the packages are brought together and compiled using an old and slightly convoluted configure/make system. This is caused mainly by the fact that the packages were brought together under the KULL name after being developed for their own purposes and with their own make systems. The need was seen for a more unified and easier-to-use make system for the entire set of KULL packages.

It was decided that an interactive graphical user interface (GUI) make system would provide the greatest benefit to developers, and ideas were presented as to features to be included in the new system (e.g., cross-

checking of package selection, inclusion of library, optimization, and compile flag options).

We began by producing a piece of code that performed the “bare-bones” tasks of allowing the user to set the options needed in a graphical environment and to integrate that with the current configure and make system. That code will be modified and augmented in order to create the final product and allow for additional features later in time. The first program was completed, integrated with the current system, and is ready for use.

Current developers of the KULL code will take over the modification and development of this program, working with it and changing it to fit their needs and the needs of KULL users.

Wavelets in Pattern Recognition

Rachel Karchin

University of California, Santa Cruz

Research Summary

I spent the summer researching the use wavelets to decompose images to extract features for pattern recognition. The results are contained in the survey paper “Classifying Images with Wavelet-based Feature Extraction,” which will appear as a UC-technical report. The paper includes an introduction to wavelet analysis, wavelet-based classifying systems, and describes in some detail the techniques used to implement such systems, a variety of applications, and results. It contains many pictorial examples to make the concepts quickly accessible, and pointers to further information on the subject.

A number of the techniques described in this paper are new. I added them to the Image Processing Requirements list that the Project Sapphire group is using as a basis for software design.

I also wrote Matlab scripts to experiment with several wavelet-based feature extraction techniques. These include calculating the energy, entropy, high frequency/low frequency energy ratio, and autocorrelation of different levels of a wavelet decomposition of various signals, and producing figures of 1D and 2D images at different resolution levels. The results appear in the survey paper.

Preconditioners for Non-Matching Grids in 3D Elasticity

Chisup Kim

Texas A&M University

Research Summary

My research is on the construction of preconditioners in the case where grids on subdomains have been coarsened independently. This coarsening strategy is a common practice in algebraic multigrid applications, which is one of the main projects at CASC, and leads to subdomain interfaces with non-matching grids. Then, to retain the approximation properties on these interfaces, the mortar finite element method, a non-conforming finite element method, is applied.

In the 3D linear elasticity application (away from the incompressible limit), such an interface is a 2D domain, where the mortar finite element method has not been thoroughly analyzed and applied. Joseph Pasciak, Raytcho Lazarov, who were visiting here from Texas A&M University, Panayot Vassilevski at CASC, and I developed a stable and efficient algorithm that can handle such 2D interfaces. A code for this algorithm is being developed and this result is expected to be published in an academic journal.

Data Mining in Astronomical Data

Imelda Kirby

University of Washington

Research Summary

Data collection threatens to overwhelm our capacity to organize and review it for scientific information and inference. Project Sapphire in the Center for Applied Scientific Computing at LLNL is developing new techniques for analyzing massive quantities of data to extract the features that are useful for scientists engaged in the research.

Our project in data mining for astronomical data began with the installation of IRAF, and other software for astronomical images. I learned how to handle the MACHO (Massive Compact Halo Object) images, and to research different automated ways of detecting asteroids in the terabytes of data MACHO contains. I also researched Self-Organizing Map (SOM) algorithms to see if they could be applied to extract features from large astronomical databases.

Software Infrastructure for Large-Scale Partial Differential Equation Simulation Codes

Matthew Knepley

Purdue University

Research Summary

My work focused on providing software infrastructure for large-scale partial differential equation (PDE) simulation codes. I became involved in the Equation Solvers Interface (ESI) effort to develop an interface standard for linear and nonlinear solvers, which led to a report co-authored with Andy Cleary. This report outlines a proposal for a framework based upon an Operator interface, which encapsulates the behavior of an arbitrary finite dimensional operator.

I have also prepared a draft proposal for Frameworks, which examines the object-oriented design of simulation software for mesh-based PDEs, as well as portability and interoperability issues. Furthermore, I looked at solving the nonlinear system that arises from fluid flow problems, using a novel projection scheme related to traditional Krylov methods.

Toward a Parallel Three-Dimensional Ideal Magnetohydrodynamics Solver

Joseph Koning

University of California, Davis

Research Summary

The simulation of ideal magnetohydrodynamics (MHD) requires the solution of a fluid dynamics problem coupled with electrodynamics. An interesting problem in this field is a compact toroid moving through a plasma contained within a tokamak fusion device. The idealization of this problem is a superconducting sphere moving through an ideal plasma. The sphere's motion will excite waves in the plasma. The energy loss through the interaction of the sphere with the plasma and magnetic field will result in the sphere slowing and eventually stopping. This problem has been treated analytically by Newcomb. The superconducting sphere in an ideal plasma serves as a proof-of-concept for the problem of a compact toroid interacting with a tokamak plasma. This problem is three-dimensional and contains fluid phenomena such as shocks and electrodynamic phenomena such as Alfvén and magnetosonic waves.

I will use an Arbitrary Lagrangian Eulerian (ALE) method couple; the ALE method to treat the fluid dynamics of the plasma; and the VFE method to treat the electric and magnetic fields. The ALE method has the ability to track interfaces; in the simulation, the sphere moves through the plasma making front tracking necessary. The VFE was chosen for its ability to treat the electric and magnetic field's continuity as it is physically. Vector basis functions

maintain normal continuity for the magnetic field and tangential continuity for the electric field.

I have been researching the abilities of the ALE methods available for use as a basis for the MHD simulation. The fluid dynamics framework I chose is KULL, an object-oriented program with an interpreted interface with the ability to treat a fluid using pure Lagrangian, free Lagrangian, Eulerian or ALE. Currently KULL does not have finite element capability. In the latter part of this year I developed an object-oriented VFE framework which can create the Hilbert and Gramm matrices for the four types of basis functions (cell, node, edge and face) as well as mapping matrices for improved performance. This framework will then be incorporated with KULL to complete the framework for the MHD simulation. I have also researched parallel linear algebra libraries and have become proficient in the use of PETSc, the portable extensible toolkit for scientific computing.

In the future, I will be coupling the VFE framework and KULL. This will entail new methods for manipulating the vector basis functions and advecting the magnetic field. Once I complete the proof-of-concept simulation, the simulation of the real physical system will require further extension of the VFE and KULL frameworks.

Superresolution of Buried Objects in Layered Media by Near-Field Electromagnetic Imaging

Sean K. Lehman

University of California, Davis

Research Summary

In many tomographic wave (either electromagnetic or acoustic) imaging systems, wavelengths and distances are such that the object under evaluation lies in the near-field. Near-field scattering includes non-radiating or evanescent fields. Evanescent fields carry high spatial frequency information. Thus, if they can be incorporated into tomographic reconstruction algorithms, higher spatial resolutions can be achieved than those predicted by the classical (Rayleigh) resolution limit. This is known as “superresolution.” Most current tomographic reconstruction methods neglect the evanescent fields.

I have been developing a physics-based theory (as opposed to an engineering-motivated approach) that seeks to understand how the evanescent fields are scattered and transformed when measured. Analysis shows that the spectrum of the evanescent fields shifts into propagating regions during the scattering and measurement process. This has never been considered previously. This “smearing” of spectral energy would be considered as noise in the current techniques which do not account for the evanescent fields.

My analysis is in a planar geometry. That is, a planar measurement system. After the dissertation, I would like to develop it for an annular or circular geometry.

Benchmarks and Models for Time-Dependent Grey Radiation Transport with Material Temperature in Binary Stochastic Media

David Miller

University of California, Davis

Research Summary

We present benchmark calculations for radiation transport coupled to a material temperature equation in binary random media. The mixing statistics are taken to be homogeneous Markov statistics where the material chunk sizes are described by Poisson distribution functions. The material opacities are taken to be constant. Benchmark values for time evolution of the ensemble average values of material temperature, energy density, and radiation transmission are computed via a Monte Carlo type method. These benchmarks are used as a basis for comparison with three other approximate methods of solution.

One of these approximate methods is simple atomic mix, which is seen to consistently over-absorb resulting

in lower steady state radiation transmission and material temperature. The second approximate model is an adaptation of what is commonly called the Levermore–Pomraning model, which we refer to as the standard model. It is shown to consistently under-absorb resulting in higher steady-state radiation transmission material temperature. We show that recasting the temperature coupling as a type of effective scattering can be useful in formulating the third approximate model, an adaptation of a model due to Su and Pomraning that attempts to account for the effects of scattering in a stochastic context. We show this last adaptation demonstrates consistent improvement over both the atomic mix and standard models.

Non-Uniform Load Balancing in SAMRAI

Scott Morris

University of Utah

Research Summary

The end product of the CSAFE project at the University of Utah is a PSE (Problem Solving Environment) that couples MPM (Material Point Method) simulating container dynamics with a continuum firespread code. Both of these pieces sit upon the SAMRAI adaptive mesh refinement infrastructure. The incorporation of the MPM code with SAMRAI required the creation of a new particle data type. These particles are free to move about the mesh and as such non-uniform distributions of the computational workload result across the mesh. Since the CSAFE project is part of the ASCI initiative, the goal is to achieve a PSE that is scalable to large numbers of processors. To facilitate this goal, SAMRAI requires a non-uniform load balancer. The objective of my work was to research an algorithm for non-uniform load balancing within an AMR framework and to implement a general purpose load balancer in SAMRAI.

The traditional approach to dividing work among processors on a grid calculation is to use recursive bisection. This approach is useful when given a large,

single domain. With adaptive mesh refinement, however, we are given the task of balancing multiple levels of refinement and irregular domains. Instead, we use a modified form of recursive bisection that can partition a collection of boxes into any number of boxes of a specified size and produce a smaller leftover box as well. This is a heuristic for chopping boxes that prepares for the next step of the algorithm: a bin-packing stage in which chopped boxes are assigned to processors. In practice, this heuristic has proved to have merit. For the linear advection sample SAMRAI code we have achieved average estimated load balances in the 90–100% range.

Further testing of the algorithm is required as different applications using SAMRAI have different needs. Specifically we will be testing the non-uniform load balancer with the MPM code currently under development at the Utah CSAFE site. Other refinements to the algorithm include consideration of communication costs and a more efficient bin-packing algorithm.

Parallel Sparse Matrix–Vector Multiplication

Jason Perry

University of Kentucky

Research Summary

Multiplying a sparse matrix by a vector is a basic task in nearly every scientific simulation code run on parallel computers. My task was to attempt to speed up this operation by adding threads to an already parallel code. The proposition rests on the dual hardware structure found in supercomputers and clusters of networked workstations. Each node of the machine (or cluster) is a complete separate computer, and messages can be exchanged between nodes to solve larger problems. Within each of these nodes is also contained a number of processors that share the memory within the node. Preexisting codes perform matrix–vector multiply in parallel by essentially dividing the rows of the matrix up among the nodes. These codes use only message passing, and therefore either ignore all but one processor within each node or pack multiple processes within one node, which causes unneeded communication overhead. Multiple threads of execution within a shared memory space allow work to be divided between processors in a node, without the need of any communication within the node itself. A program code that combines message passing and threading should reduce the amount of inter-process communication on parallel computers, thereby speeding up the computation.

My first step in this experiment was to write code that would allow me to input matrices for testing. I wrote a program in C to read matrices stored in the standard Harwell–Boeing format. I then proceeded to write from scratch a parallel matrix–vector multiplication code that

uses the MPI message passing standard and no threads. This code was tested on both the IBM Pacific Blue supercomputer and a cluster of DEC Alpha workstations. I tested and timed this code and refined it until I felt it was sufficiently fast, as well as representative of existing codes. Then, I added code that uses the Pthreads threading standard to thread the multiplication within each node. I created several new versions of the program, each with a different method of threading the loop that does the multiplication. Some methods use dynamic load balancing to evenly distribute work among the threads, while others simply divide up the rows of the matrix beforehand. Some versions perform the multiplication and communication between nodes at the same time, by assigning different tasks to different threads. These were all extensively tested with a variety of matrices from scientific applications, as well as specially generated regular matrices with up to 300 million entries. The goal was that my threaded code would run faster than a non-threaded code running the same problem on the same number of processors.

The results have been positive, showing that threading can provide a modest speedup if done correctly. Also, important insights have been gained into some of the performance limitations of new parallel machines. I have prepared extensive documentation, which will be further revised, and which I hope to publish, and have given two presentations on this work.

Scalable Domain Decomposition Algorithms for Resolving Contact Surfaces in ALE Computations

Tim Pierce

University of California, Davis

Research Summary

In the Lagrangian approach to modeling material media, the overall volume of material under consideration is subdivided into a discrete set of zones, together with their connecting faces, edges, and nodes: i.e., a mesh. The connectivity of the mesh determines which nodes belong to each zone, from which can be deduced which zones contain a given node. This information establishes the proximity relations of all zones.

As the material moves in time under the influence of internal stresses, body forces, and boundary conditions, the mesh follows the material, so that a given zone always contains the same material. For an explicit time advancement scheme, the behavior at a given node or zone over a single timestep is dictated only by material in its own and neighboring zones, which can be determined from the mesh connectivity. The timestep is chosen, with reference to the speed at which signals can propagate through the material (typically the sound speed), so that this locality of influence is a valid assumption.

If the model includes more than one discrete body, the interaction of the bodies must also be considered, and mesh connectivity is no longer sufficient to determine proximity. The bodies may come together or separate, or may slide along each other, each exerting a boundary force on the other. Though the connectivity of the mesh representing each individual body is constant, which zones are adjacent across the contact surface is entirely dynamic and unpredictable.

A standard approach to parallelizing a Lagrangian dynamics code involves dividing the mesh into a number of submeshes, or subdomains, of approximately equal size, and assigning the subdomains to separate processors. A good decomposition minimizes communication, as only neighbor information is required. Furthermore, since the connectivity is constant in time the communication pattern

can be set up during an initialization phase, and only the variable data itself updated each cycle.

This simple domain decomposition method of parallelization breaks down at contact surfaces, or at least an additional mechanism must be supplied to a) determine proximity relations across contact surfaces, b) distribute the calculation of the contact forces across the parallel machine, and c) gather the necessary data from across the machine together on the processor where that part of the surface is to be calculated.

The task of calculating contact forces can be divided into two steps, contact detection and contact enforcement. First, one must determine if and where two surfaces are in contact. Then, one can apply suitable balancing forces at those points.

In a serial environment, the detection problem can always be solved by brute force. The position of each node on one side can be compared to the positions of all faces on the other side, the closest face determined, and then penetration checked. In a parallel environment, however, even the brute force approach is not available, since the closest face on the other side may be on a different processor. As time changes, the closest face and perhaps the processor assigned to that face may also change.

As opposed to the relatively simple problem of parallelizing a Lagrangian dynamics code on a single mesh, the parallelization of contact surfaces provides formidable challenges, particularly if the task must be done in a scalable fashion as the mesh size and processor count are increased proportionately.

The purpose of this research is to develop a particular algorithm designed to provide a truly scalable solution, to implement it in a major dynamics code (ALE3D), and to demonstrate scalability well beyond 1,000 processors.

ROAM Using Surface Triangle Clusters

Alex Pomeranz

University of California, Davis

Research Summary

As terascale datasets become commonplace, new methods are needed to visualize these enormous amounts of data accurately and quickly. Current display techniques are CPU-heavy and do not come close to utilizing the full potential of the graphics cards in the consumer market. Our accelerated display technique builds upon the view-dependent optimization technique presented in ROAM (Duchaineau) with the intention of allowing the user to view arbitrary 3D surfaces in real-time.

The original ROAM paper used a BinTree hierarchy to store various approximation levels of our 2D mesh/terrain. By taking some cut across this tree, such that a certain number of polygons is displayed or a certain maximum error threshold is met, we can optimally display our mesh given the constraints of our error metric in real time. Graphics cards today are capable of rendering 9 million shaded and textured triangles per second; however, the original ROAM paper sent only 6,000 triangles/frame at 30 frames/second, for a total of 180,000 triangles. This amounts to a 2% utilization of our graphics hardware.

In our new method, we take a more coarse-grained approach. After the cut across the BinTree hierarchy is found, rather than display individual triangles, we display triangle clusters, already precomputed and pre-packaged for quick output. Assuming 256 triangles/cluster, we would then be pushing 180,000 triangle clusters at 256 triangles/cluster, or 4.6 million triangles/second. This not only gives a much better approximation of our actual mesh, but it also offers a 25 times increase in the utilization of our graphics hardware.

We intend to adapt this method to 3D meshes. This gives our algorithm a much wider base of potential users. Currently, we are working on the initial prototype for this system. We already have a large-object oriented framework up and running, and are incrementally adding functionality to it. Some issues to be addressed are the adaptation of the algorithm to 3D meshes, “prepackaging” triangle clusters, and “paging clusters to disk” issues.

A Multigrid Strategy for Accelerating Steady-State Computations of Waves Propagating with Curvature-Dependent Speeds

Jonathan Rochez

University of California, Davis

Research Summary

The need for a fast and accurate representation of lighting times of detonation of materials has arisen in several problems. Such a representation is called a burn table and is simply an array of times at which the detonation burn front crosses a particular point relative to the detonation point(s). In a high-explosive material, the creation of a burn table will allow the elimination of solving chemical reaction ordinary differential equations (ODEs) and feed in source terms to the reactive flow equations for solution of the system of ignition of the high explosive material. Standard iterative methods show a quick reduction of the residual followed by a slow final convergence to the solution at high iterations. Such systems are excellent choices for the use of multigrid methods to speed up convergence, even on a nonlinear system such as this. During the past year, we have attempted to develop this code in 2D and 3D using two differencing schemes for the multigrid method, and numerical steady-state solutions to the eikonal equation on a rectangular grid were conducted.

Numerically, two different approaches have been investigated with the multigrid method. The first approach is organizing the calculation point by point on the grid using a finite differencing scheme. The advantages of this approach include its speed in calculation time and ease of implementation, while its downfall is having to remain on a regular orthogonal grid. During this past year, code was developed to solve the eikonal equation on a unit square geometry with one or several detonation points involved. Comparisons of several sized problems with several multigrid V-cycles have been compared. In all cases, the speed was significantly faster for the

multigrid method versus a regular calculation on a single grid. Also, the larger the problem size became, the faster the speed-up that was observed. The best experimental results as seen in the comparison table below show a factor of 60 reduction in wall clock time.

The second approach is one of using a finite element scheme in a zone by zone organized calculation. Advantages include unstructured grids for the formation of non-regular geometries and fewer data points, while the disadvantages include larger calculation times and being somewhat more difficult to implement. Code has been developed to test this method as well. No results have been obtained yet, due to the unstable and non-converging behavior of using a linear “hat-function” basis set.

Current plans for improvement include expanding the geometries for the finite differencing scheme with the incorporation of curvature in the calculation of the burn table. Error analysis and comparison to physical data must be performed as well. The finite element approach will be fully implemented, while much work is still needed on the grid generation section—specifically prolongation and relaxation between grid levels.

Comparison of single grid versus multigrid for select problems

Problem	Calculation time (sec)
2D, 128x128, single grid	1889.33
2D, 128x128, multigrid	314.29
3D, 33x33, single grid	66723.40
3D, 33x33, multigrid	1023.49

Gamma-Ray Bursts: Numerical Modeling of Electron–Positron Pair Plasmas in Cataclysmic Astrophysical Phenomena

Jay Salmonson

University of California, Davis

Research Summary

Despite three decades of intense scientific scrutiny, gamma-ray bursts have remained one of astronomy's biggest unsolved mysteries. Recent observational breakthroughs have allowed us to learn much about these big, brief, brilliant bangs seen from across the cosmos, but their origin remains a mystery. In this work, we study three progenitor models: a neutron star binary system near its last stable orbit, a charged black hole, and the collapse of a globular star cluster. All of these scenarios result in a common theme, the relativistic expansion of a super-heated electron–positron–photon plasma. Thus, we study the evolution of, and emission from, this plasma as it might result from these three progenitors using numerical general relativistic hydrodynamic simulations. This emission is then compared with that of gamma-ray bursts to test the feasibility of each of these models as a gamma-ray burst progenitor.

Implementing the Full Approximation Scheme Algorithm

Kevin Scully

University of California, San Diego

Research Summary

We have created an implementation in C of the FAS (Full Approximation Scheme) multigrid algorithm developed by Achi Brandt to apply to discretizations of nonlinear partial differential equations (PDEs). This algorithm may permit the solving of nonlinear PDEs with speed and efficiency comparable to that of the linear multigrid method on linear PDEs. While not much theory exists to support the effectiveness of this algorithm, its power has been observed in certain implementations.

Under the guidance of Carol Woodward, an FAS algorithm using a Point Newton finite difference weighted Jacobi smoother, semicoarsening to generate coarse grids, full-weighting and injection for restriction and linear interpolation, was developed to handle linear and nonlinear problems in up to three dimensions, where the domain is of tensor-product type. This code reproduces results of benchmark Matlab FAS codes written by Van Henson on elliptic linear and nonlinear Dirichlet boundary problems.

This code uses many data structures and functions to manipulate these data structures from the *hypr* library, developed at LLNL. In particular, it makes liberal use of the “hyprebox” data structure and has been written to handle multiple box domain partitions with an eye to a parallel computing implementation. Many of the routines dealing with boxes (such as the restriction and interpolation functions) were patterned after comparable routines in the linear multigrid solver PFMG, also developed here at LLNL by Rob Falgout and others.

I plan to explore the possibility of interfacing this code with finite element discretizations of Einstein’s equations produced by Michael Holst’s MC package. To use this algorithm in space-time finite element discretizations (where the finite elements have a time component), I will have to extend this algorithm to four dimensions.

Flexible GMRES in CVODE and PVODE

Alexandru Tamasan

University of Washington

Research Summary

CVODE is a nonlinear solver for ordinary differential equation initial value problems developed by A. Hindmarsh and S. D. Cohen at LLNL. PVODE is a parallel version. The algorithm executes an outer time loop with an adaptive time step. Newton's method is applied in solving an implicit time discretization scheme (e.g., backward differentiation or Adams) on each time step. A linear solver is required for the corresponding Jacobian equation. An iterative method based on the GMRES algorithm is available, called CVSPGMR. In 1993, Saad noticed that for the price of a little extra storage, GMRES accommodates a variable right preconditioning. This new linear solver is called Flexible GMRES. My task was to implement it within the CVODE package, resulting in a module called CVSPFGMR. Despite the fact the solution is formed differently, the residual solves the same least squares problem as in the GMRES algorithm. Hence we can use the same stopping criteria. Moreover, the conver-

gence results for the Inexact Newton–GMRES carry over to the Inexact Newton–FGMRES algorithm.

In order to compare the two methods, a test problem previously run with CVSPGMR was run with CVSPFGMR. A variable preconditioning on the right was realized by requiring a fixed tolerance on the number of Gauss–Seidel preconditioner iterations, rather than a fixed number of iterations. For similar performance in terms of measured outputs, the numbers of iterations could be reduced by a factor of more than two, making a variable preconditioner a better choice. The disadvantage is the increase in the memory space.

Work in progress deals with the mixed initial-boundary value problems for a 3D heat equation. PVODE runs with CVSPFGMR as the linear solver, CG with SMG from the *hypr* package as the preconditioner, combining two independently developed CASC codes.

Unsplit ALE Using an Augmented Lagrangian Mesh Velocity

Jay F. Thomas

University of California, Davis

Research Summary

The Arbitrary Lagrangian–Eulerian (ALE) method has been successfully applied to the simulation of variety of problems including impact phenomena, material processing (e.g. metalforming) and fluid–structure interaction. The ALE method combines the Lagrangian approach, in which the computational mesh moves with the material, with the Eulerian approach wherein the mesh is fixed and the material flows through it. The Lagrangian approach simplifies application of boundary conditions and tracking of material interfaces, but the mesh can become severely distorted or entangled—eventually halting the computation.

The Eulerian approach allows arbitrarily large deformations and turbulent flow but treatment of boundary conditions and material interfaces may require special accommodation. The ALE formulation attempts to circumvent these difficulties and obtain the advantages of both methods by allowing the mesh to be moved in an arbitrary manner. The issue becomes one of selecting the mesh motion so as to control grid distortion and or solution error.

ALE methods have typically been implemented as uncoupled or “split” procedures consisting of a Lagrangian phase followed by an Eulerian or advection phase in which the mesh velocity is calculated (remeshing) and the state variables are remapped to the updated mesh positions. There is no time step associated with the advection phase. The advection step can add considerable computational expense, particularly if frequent remeshings are needed. Many ALE codes adopt a

mesh relaxation procedure to move the nodes with an eye towards controlling grid distortion. Mesh relaxation is based upon elliptic mesh generation methods. The coordinates of the nodes are formulated as the solution to an elliptic problem and Jacobi iteration is used instead of a direct solution method. This results in a stencil that can be applied efficiently over the mesh. Only a few sweeps of Jacobi are used at each time step.

The goal of the effort is to develop and investigate unsplit ALE formulations and associated computational procedures with particular attention to control of mesh distortion. With an unsplit formulation, there is no operator split and hence no separate Lagrangian and Eulerian phases. In principle, greater accuracy is possible in comparison to the split procedure. Controls for mesh distortion are specified directly as part of the differential equation for the mesh velocity, as opposed to indirectly via mesh relaxation. The mesh may thus be considered to deform, and potentially adapt to the solution, continuously. Initial effort is focused on mesh velocities of the form: $u + f(t, x)$ where u is the material velocity, x is spatial position and t is time. f is a “correction” term to prevent mesh distortion and entanglement. Preliminary experiments performed using the inviscid Burgers’ equation indicate promise for the method.

Future plans are to develop the formulation for a 1D system of equations to aid in choosing appropriate candidates for the $f(u, x, t)$ term. The technique would then be extended to 2D.

Sparse Multiple Semicoarsened Grids

Heidi K. Thornquist

Rice University

Research Summary

Much work has been done developing robust multigrid methods for solving systems arising from finite difference, finite volume or finite element discretizations of the diffusion equation, in the presence of discontinuous coefficients and anisotropy. However, all of these methods have higher computational and memory costs than methods based on point smoothers and simple local grid transfer operators. The goal of this work is to develop an algorithm that is both robust and efficient.

The new method we developed is a modification of the multiple semicoarsened grids (MSG) method of Naik and Van Rosendale. The robustness of this method was studied by Oosterlee and Wessling and a new, more efficient variant of this algorithm was recently proposed by Washio and Oosterlee. Our approach for improving efficiency is based on the relationship between the MSG algorithm, the hierarchical-basis multigrid method, and the sparse grid methods of Zenger. Our new method, denoted SparseMSG, has been tested on a variety of 2D and 3D scalar diffusion problems, where, in iteration count, it beats LLNL's PFMG. In the future, we intend to optimize SparseMSG for parallel implementation.

Selecting Data Mining Tools

Ben Tobin

Northern Arizona University

Research Summary

Scientific data that has been carefully collected over time can now be exploited by new data mining technology. For use in anticipated LLNL terabyte applications, a data mining tool must be able to handle hundreds of data columns, handle flat files, and be parallel. To evaluate data mining options for LLNL, I began with the *Two Crows* report, a 370-page report evaluating the 25 largest data-mining products on the market, of which 8 support parallel architectures. In consultation with CASC scientists, we have narrowed the field to five data mining products: Intelligent Miner, SPSS, CART, Darwin, and Clementine.

We presented the results of our trials and the ease of interaction with these various products. It is possible, however, that the data-mining tool that we need has not yet been developed.

Selected references

Herbert Edelstein, *Two Crows Data Mining '99 Technology Report*, Two Crows Corporation, Potomac 1999.

JSL Decision Systems Inc. Clementine
<http://www.isldsi.com>

SPSS
<http://www.spss.com/datamine>

Salford Systems, CART
<http://www.salford-svstems.com>

Neo Vista Software, Decision Series
<http://www.hp.com/fsi/partners/partner585.html>

Silicon Graphics, MineSet
<http://www.sgi.com/software/mineset>

Thinking Machines Corporation, Darwin
<http://www.think.com>

Domain Decomposition on Non-Matching Grids

Stanimire Tomov

Texas A&M University

Research Summary

I worked with Raytcho Lazarov and Panayot Vassilevski on two articles, “Interior Penalty Discontinuous Approximations of Elliptic Problems” and “Discontinuous Approximations of Advection-Reaction Equations”. The first article studies a penalty discretization of second-order elliptic boundary value problems, which is first-order accurate if one uses piecewise linear elements. The second one proposes a conservative (finite volume) approximation of advection-reaction problems. I finished the implementation (in C++) for both methods and provided extensive numerical results. I worked on *a posteriori* error estimators and adaptive refinement and applied some local refinement techniques for the penalty formulation for diffusion problems. Work is now in progress is work on mortar approx-

imation for convection-dominated diffusion problems (with upwinding), and derivation of local error estimators for purely convection-reaction problems.

A significant part of my work was on 3D mesh generators, design and implementation in C++ of data structures for multilevel adaptive grid refinement, implementation of various local error estimators, and tests of their efficiency. The experience of Mike Holst and Randy Bank, who were visiting CASC at the same time and with whom I had discussions, turned out to be very helpful. For the 3D mesh generation I used NETGEN, a 3D mesh generator based on advancing front method, efficient in generating coarse meshes. The mesh refinement is based on Douglas Arnold’s bisection algorithm.

Interface Design for Visualization Applications Using Adaptive Mesh Refinement

Kevin Vlack

University of Illinois, Urbana–Champaign

Research Summary

Adaptive Mesh Refinement (AMR) has recently become a popular method for focusing computational resources on subregions of a domain relevant to the stability and accuracy of computational experiments. Subregions created through AMR may then be re-assembled to represent the entire domain, but the complexity of this reassembly complicates numerical analysis, and visualization tools have become necessary to study the results of these experiments. Such tools require efficient methods of data representation which can be easily interpreted by a human researcher.

AMRVIS, developed by Vince Beckner of Lawrence Berkeley National Laboratory, is an application used to visualize data from files produced by *hypr*, a library for AMR solvers for problems in fluid dynamics and radiation transport. Our new version provides an interface improvement to its previous release, reducing the clutter in the GUI and introducing menu item accelerators for every option in the tool, such as viewing datasets and

selecting different levels of refinement for view. A module has been added for viewing older versions of the data files produced by the *hypr* code, which can now be viewed in parallel. A new 1D line plotting utility allows viewing the values along a horizontal or vertical line of the subdivided computational domain on a traditional Cartesian plot. This feature also supports animation and level-switching. Such capabilities are important to the end-user and the *hypr* programmer for purposes of software evaluation and for understanding the physical processes behind experimental results.

The *hypr* code is currently under development for parallel 3D simulations for fluid dynamics and diffusive radiation transport. As the code nears completion, the code for AMRVIS must be able to interactively support the visualization of selections of a 3D dataset. While the current software and its new version allows viewing in 3D, its methods in data selection and visualization are still rather simple and must undergo further development.

Data Locality Optimization for the ROSE C++ Preprocessor

Christian Weiss

Technische Universität

Research Summary

The ROSE preprocessor is part of LLNL's Overture project for solving systems of partial differential equations on general multipatch grids. I implemented a two-dimensional red-black Gauss-Seidel relaxation method using the A++ class library to get familiar with the usage and implementation of ROSE. ROSE was not able to transform codes which contains multidimensional arrays. Hence, I modified ROSE so that it is now able to correctly handle and transform C++ code which uses A++ arrays with more than one dimension.

To automatically apply data locality transformations with a preprocessor, several features are needed within the code transformation tool. I identified additional features I would need within ROSE to implement data locality optimizations. Array padding was selected, since it requires the least features to be added to the ROSE preprocessor. Array padding is a technique that changes the memory layout of arrays to reduce the number of conflict misses in an application. The library I implemented provides array padding heuristics as well as tools to apply padding to array descriptions. The library can be used to implement a code transformation within ROSE as well as to implement runtime support for A++ array constructors to improve the memory layout of dynamically allocated arrays.

A Reduced Grid Method for a Parallel Global Ocean General Circulation Model

Michael Everett Wickett

University of California, Davis

Research Summary

A limitation of many explicit finite-difference global climate models is the timestep restriction caused by the decrease in cell size associated with the convergence of meridians near the poles. A computational grid in which the number of cells in the longitudinal direction is reduced toward high-latitudes, keeping the longitudinal width of the resulting cells as uniform as possible and increasing the allowable timestep, is applied to a three-dimensional primitive equation ocean-climate model.

This “reduced” grid consists of subgrids which interact at interfaces along their northern and southern boundaries, where the resolution changes by a factor of three. Algorithms are developed to extend the finite difference techniques to this interface, focusing on the conservation required to perform long time integrations, while preserving the staggered spatial arrangement of variables and the numerics used on subgrids. The reduced grid eliminates the common alternative of filtering high-frequency modes from the solution at high-latitudes to allow a larger timestep and reduces execution time per model step by roughly 20%. The reduced grid model is

implemented for parallel computer architectures with two-dimensional domain decomposition and message passing, with speedup results comparable to those of the original model. Both idealized and realistic model runs are presented to show the effect of the interface numerics on the model solution.

First, a rectangular, mid-latitude, flat-bottomed basin with vertical walls at the boundaries is driven only by surface wind stress to compare three resolutions of the standard grid to reduced grid cases which use various interface conditions. Next, a similar basin with wind stress, heat, and fresh water forcing is used to compare the results of a reduced grid with those of a standard grid result while exercising the full set of model equations. Finally, global model runs, with topography, forcing, and physical parameters similar to those used for ocean-climate studies, are advanced to a near equilibrium state for both the reduced grid and the standard grid. Differences between the two are presented for typical fields of interest, and very little degradation of the solution due to the reduced grid is observed.



Institute for Scientific Computing Research

Scalable Linear Solvers Workshop Report



Working to Solve Large-Scale Linear Systems

Robert D. Falgout and Jinchao Xu

Synopsis of Workshop Events

The multigrid method has been actively investigated by researchers for many years, but how can it be made more practical in applications? This question was discussed in detail at a workshop sponsored by CASC, the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory.

This three-day (June 23–25, 1999) Oberwolfach-style workshop was held in the beautiful Wente Vineyards of Livermore, California. Just a few minutes' drive from CASC's office, the informal setting was ideal for the size and format of this workshop. About 40 invited participants attended (mostly from the United States, with a few from Europe) and joined in lively discussions during and after the workshop sessions.

This workshop was one aspect of the many research outreach activities sponsored by CASC and the Institute for Scientific Computing Research (ISCR). Under the direction of Steven Ashby, CASC has grown from 12 scientists in 1996 to more than 80 at present. CASC conducts collaborative scientific investigations that require the power of high-performance computers and the efficiency of modern computational methods. Its research and development activities are applications-driven, and focus on those LLNL programmatic objectives requiring advanced computational technologies. The Center's core competencies include high-performance computing, computational physics, numerical mathematics, and computer science.

The workshop was organized by researchers in the Scalable Linear Solvers project in CASC (http://www.llnl.gov/CASC/linear_solvers/). The project develops scalable algorithms and software for solving large, sparse linear systems of equations on parallel computers. A major research focus is multigrid, one of

the most powerful solution methodologies for solving large-scale systems. Ashby and his colleagues in CASC have placed the multigrid method at the center of their algorithmic research development. They intend to accelerate and broaden its application, which will have a notable impact on scientific simulation.

The multigrid method has been shown to be extremely efficient both in theory and in practice. Traditionally, however, it has been somewhat problem-dependent and difficult to use; particularly in a software library setting. In recent years, there has been a desire to make the method more accessible to general users. This has led to a resurgence of interest in a class of algorithms originally developed in the early Eighties, algebraic multigrid methods.

Similar to classical multigrid methods, algebraic multigrid methods exhibit optimal computational complexities for many applications, but are relatively much easier to use (as are other algebraic algorithms such as Gaussian elimination). These algorithms were the major focus of the workshop. In CASC, a large group of researchers works on this subject, including Andy Cleary, Robert Falgout, Van Henson, Jim Jones, Barry Lee, Beth Ong, Charles Tong, Panayot Vassilevski, and Ulrike Yang. The group has formed several collaborations, including a close partnership with University of Colorado at Boulder researchers Tom Manteuffel, Steve McCormick, and John Ruge. As the participants are mostly very active researchers on the subject, the atmosphere of the workshop was very lively and down-to-earth.

The workshop began at 9 a.m. on June 23, 1999, with Ashby's opening remarks. Falgout then explained the meeting format and program, one that encouraged an atmosphere of open exchange. Each morning and after-

Working to Solve Large-Scale Linear Systems (continued)

noon session was devoted to a special topic. Sessions opened with a few short presentations to enumerate the issues, then lively discussions followed.

The first session of the workshop, organized by Raytcho Lazarov, related to domain decomposition techniques, and focused primarily on methods for non-matching grids. Three sessions were devoted to algebraic multigrid methods. Organized by Jim Jones and Panayot Vassilevski from LLNL, two of the sessions centered around recent algorithmic advances for improving robustness. Van Henson from LLNL organized the third session, which covered the important and difficult problem of developing parallel algebraic multigrid algorithms. Another session, organized by Falgout, discussed issues related to implementing and designing multigrid methods for today's high-performance computers, where parallelism and cache-performance play a dominant role. Finally, Edmond Chow from LLNL organized a session on multilevel ILU, a method that has much in common with algebraic multigrid, but one that brings a different algorithmic perspective to the table.

In addition to the special sessions, another major activity was the afternoon group discussions. During the lunch breaks, participants could either stick together to form their own small work groups or join an organized discussion. As it turned out, most participants joined the organized sessions, which were meant to be informal and deliberately provocative.

The first such discussion was organized by Jinchao Xu, and asked the question, "Which problems can and cannot be

solved with multigrid, and why is multigrid not commonly used in practice?" As expected, this generated much controversy. In fact, the question itself was the first thing to be debated, as some of the participants agreed with the statement that multigrid was not commonly used in practice, and others disagreed. Xu constructed a web page (using his laptop) with a list of important problems, and solicited input on those problems for which multigrid is an effective solver. Again, much debate ensued, and although the details could not quite be agreed upon, an overall consensus emerged that multigrid methods can be effective solvers for a remarkably wide range of problems. This discussion carried over to the next lunch.

The second afternoon discussion occurred on the third day of the meeting and was organized by Tom Manteuffel. The similarities and differences between multigrid, domain decomposition, and aggregation methods were explored, and the question was asked, "Can one establish a common framework into which these all fit?" Manteuffel outlined a framework in which methods are viewed as "approximate" Schur complement methods. This also generated much healthy discussion, and a relevant early framework based on space decomposition and subspace correction was mentioned.

On the basis of the productivity of the workshop and the enthusiasm of its participants, CASC plans to sponsor three or four workshops per year on subjects related to its core mission. For fiscal year 2000, these will include scalable nonlinear solvers, transport methods, and mining large scientific data sets. For information on future workshops, watch the webpage at <http://www.llnl.gov/casc/workshops>.